arXiv:1701.06739v1 [stat.ME] 24 Jan 2017

Partial Bridging of Vaccine Efficacy to New Populations

Alexander R. Luedtke and Peter B. Gilbert

Vaccine and Infectious Disease Division, Fred Hutchinson Cancer Research Center

Abstract

Suppose one has data from one or more completed vaccine efficacy trials and wishes to estimate the efficacy in a new setting. Often logistical or ethical considerations make running another efficacy trial impossible. Fortunately, if there is a biomarker that is the primary modifier of efficacy, then the biomarker-conditional efficacy may be identical in the completed trials and the new setting, or at least informative enough to meaningfully bound this quantity. Given a sample of this biomarker from the new population, we might hope we can bridge the results of the completed trials to estimate the vaccine efficacy in this new population. Unfortunately, even knowing the true conditional efficacy in the new population fails to identify the marginal efficacy due to the unknown conditional unvaccinated risk. We define a curve that partially identifies (lower bounds) the marginal efficacy in the new population as a function of the population's marginal unvaccinated risk, under the assumption that one can identify bounds on the conditional unvaccinated risk in the new population. Interpreting the curve only requires identifying plausible regions of the marginal unvaccinated risk in the new population. We present a nonparametric estimator of this curve and develop valid lower confidence bounds that concentrate at a parametric rate. We use vaccine terminology throughout, but the results apply to general binary interventions and bounded outcomes.

Keywords: bridging; external validity; generalizability; partial identification; transportability; vaccine efficacy

1 Introduction

Randomized clinical trials are the gold standard for evaluating the effect of a new intervention in a population. However, it is not always possible or appropriate to conduct a Phase 3 efficacy trial. In these cases, it is desirable to "bridge" results from earlier trials to a new setting that may differ in population, treatment version, or exposure levels. Being able to extrapolate efficacy signals to a new population is a special case of generalizability or external validity (Cole and Stuart, 2010), and is often referred to as transportability (Bareinboim and Pearl, 2012, 2013; Pearl and Bareinboim, 2014). Following the clinical trials literature, we will instead refer to this extrapolation as "bridging" (Gilbert and Hudgens, 2008; Gilbert et al., 2011; Fleming and Powers, 2012).

There have been several recent trials establishing the efficacy of the CYD-TDV vaccine in children. In particular, CYD14 found 56% per-protocol efficacy of the vaccine in Asian children aged 2-14 years (Capeding et al., 2014), CYD15 found 61% per-protocol efficacy the same vaccine in Latin American children aged 9-16 years (Villar et al., 2015), and CYD23 found 30% per-protocol efficacy in Thai children aged 4-11 years (Sabchareon et al., 2012). All trials featured a 2:1 vaccine:placebo randomization scheme, with respective sample sizes of approximately 10,000, 21,000, and 4,000 and respective virologically confirmed dengue event counts of 250, 397, and 134. In each trial, efficacy was shown to increase with age and prior exposure to dengue (Capeding et al., 2014; Villar et al., 2015; Sabchareon et al., 2012). While several countries have licensed the CYD-TDV vaccine for individuals from 9–45 or 9–60 (WHO, 2016), many countries have still not licensed this vaccine due to debates over the appropriate ages of indication. Nonetheless, the high estimated efficacy in the existing trials may make further Phase 3 trials unethical. There is thus compelling motivation to bridge the results of the earlier efficacy trials completed in children to the adult population.

The focus of most existing work on bridging is on bridging the results from a single efficacy trial to a new population. Unsurprisingly, some assumptions must be made to bridge between the two studies. The primary standard assumption is that (i) the two populations have equivalent outcome regressions, i.e. that the mean outcome within each treatment/covariate/immune response biomarkers stratum is constant, and (ii) that the support of the treatment/covariate/immune response biomarker distribution in the population to which the researchers wish to bridge is contained in the corresponding support in the population in which the trial was completed (Bareinboim and

Pearl, 2012, 2013; Pearl and Bareinboim, 2014). Recently, an efficient, double robust estimator was established when these assumptions hold (Rudolph and van der Laan, 2016). There has also been some work on establishing transportation formulas when there are multiple settings in which full observations are observed in more than one setting, which in our setting means that there is more than one completed efficacy trial (Lee and Honavar, 2013; Bareinboim and Pearl, 2014).

Assumption (i) is often not plausible in the infectious disease setting unless one adjusts for the level of pathogen exposure, but in most efficacy trials there is only coarse exposure data (e.g., age in dengue trials). Furthermore, adjusting for level of pathogen exposure can violate Condition (ii) in some scenarios. To overcome this problem, Gilbert and Huang (2016) posited a bridging assumption on the ratio scale, arguing that the vaccine versus placebo risk ratios in the two populations should be approximately equal within each covariate-specific principal stratum because this ratio reflects the vaccine's effect on biological susceptibility to pathogen infection or disease. Because knowing the risk ratio does not uniquely determine the outcome regression, there remains an undetermined degree of freedom. One possibility is to estimate the outcome regression among unvaccinated participants in the new population using epidemiological data (Gilbert and Huang, 2016).

There is often severe underreporting of infection or disease incidence in the available epidemiologic surveillance data. For example, recent studies comparing active and passive surveillance (by national surveillance programs) of dengue incidence have seen up to 19-fold underreporting in the passive surveillance (Nealon et al., 2016; Sarti et al., 2016). Though one could attempt to account for the underreporting by specifying an inflation factor for the passive surveillance data, there will likely be uncertainty around what exactly this inflation factor should be. For this reason, in this work we focus on a less rigid approach to account for the unknown unvaccinated risk allocation. In particular, we derive the worst-case allocation under constraints that we will impose in the next section. These constraints can be estimated via epidemiologic data or existing trial data. This then yields the lower bound on the vaccine efficacy in the new population that we wish to estimate using double robust methods (e.g., Robins and Rotnitzky, 1995; van der Laan and Robins, 2003; van der Laan and Rose, 2011). Identifying a parameter that provides a bound on an unidentifiable parameter is often referred to as partial identification (Manski, 2003), and these identification problems have received considerable attention over the last several decades (e.g., Manski, 1990; Horowitz and Manski, 2000; MacLehose et al., 2005). Because we are specifically focused on bridging a parameter from one or several populations to a new setting, we refer to this

exercise as partial bridging, though really our objective is a special case of partial identification.

Organization of article. Section 2 describes the observed data and presents a lower bound on the vaccine efficacy in the new population. Section 3 derives a first-order approximation to the lower bound of inference that will play a key role in our estimation scheme. Section 4 describes our estimation scheme, with Section 4.1 focusing on the case where the conditional vaccine efficacy surface is not known to be constant across completed efficacy trials, and Section 4.2 describing how to improve the precision of the lower confidence bound when the conditional efficacy surface is known to be the same across the completed efficacy trials. Section 5 presents a simulation study. Section 6 concludes with some brief remarks.

All of our proofs can be found in Appendix B. Appendix A presents technical regularity conditions used to establish the validity of our estimator. Appendix E extends our results to two-phase sampling designs and monotonic vaccine efficacy curves. Appendix F presents simulation results under a two-phase sampling design.

2 Problem setup

2.1 Notation and bridging assumptions

Before presenting our problem, we introduce a few basic pieces of notation. For functions f and g mapping from some space $\mathcal{X}$ to another space $\mathcal{Z}$, we write $f \equiv g$ to denote equality of $f(x)$ and $g(x)$ for all x. We also write $f \equiv 0$ to denote that $f(x) = 0$ for all x. We use "$\triangleq$" to denote a definition, e.g. $f(x) \triangleq x$ denotes the identity map. When we want to emphasize that f is a function, we refer to f as $w \mapsto f(w)$. As is standard in the empirical process literature, we let Pf denote the expectation of $f(X)$ when X is drawn from the distribution P (e.g., van der Vaart and Wellner, 1996).

Consider the data structure (W, A, Y), where W is a (possibly multivariate) baseline (pre-randomization) biomarker, A is a vaccination indicator, and Y is the outcome of interest that occurs subsequent to vaccination. While all of the results in this work hold for general bounded Y, we focus on the case that Y is an indicator of infection or disease, since this is our primary case of interest. To avoid introducing additional burdensome notation, we assume that the outcome Y is observed on all individuals, though we note in the Discussion (Section 6) that the extension to right-censored outcomes is

straightforward. Suppose that we have already observed trials in Populations $1, \ldots, S$. Each trial s consists of n_s independent and identically distributed (i.i.d.) observations $O_s[1], \ldots, O_s[n_s]$ of the structure $O_s \triangleq (W, A, Y) \sim P_s^0$, where P_s^0 is known to belong to the model $\mathcal{M}_s$ that at most places restrictions on the conditional distribution of A given W and such that each $P_s' \in \mathcal{M}_s$ satisfies $\delta < \min_{a \in \{0,1\}} P_s'(A = a|W)$ with P_s^0 probability one. Thus, we assume that all efficacy trials tested the same vaccine versus control/placebo and collected the same W and Y following the same protocol. While it is not essential that the data from each P_s^0 be drawn from a randomized trial, we refer to data from P_s^0 as data from "trial s" because that is our primary application of interest.

From each trial s we have an estimate of the vaccine efficacy, given by $1 - \frac{\mathbb{E}_s^0[\mathbb{E}_s^0[Y|A=1,W]]}{\mathbb{E}_s^0[\mathbb{E}_s^0[Y|A=0,W]]}$, where $\mathbb{E}_s^0$ is the expectation operator under P_s^0. If A is independent of W, then this reduces to estimating a covariate-unadjusted vaccine efficacy. Our objective is to estimate the vaccine efficacy in a new population $\star$. Were we to run a trial in this population, we would observe i.i.d. copies of $(W, A, Y) \sim P_\star^{0,F}$, and we could then estimate the vaccine efficacy

$$\Psi(P_\star^{0,F}) \triangleq 1 - \frac{\mathbb{E}_\star^{0,F}\left[\mathbb{E}_\star^{0,F}[Y|A=1,W]\right]}{\mathbb{E}_\star^{0,F}\left[\mathbb{E}_\star^{0,F}[Y|A=0,W]\right]},$$

where we use $\mathbb{E}_\star^{0,F}$ to denote the expectation operator under $P_\star^{0,F}$. In practice we do not observe a trial from population $\star$, but rather a size $n_\star$ i.i.d. sample of observations $O_\star[1], \ldots, O_\star[n_\star]$ containing only $O_\star \triangleq W \sim P_\star^0$, where $P_\star^0$ is the marginal distribution of W under $P_\star^{0,F}$. We denote the nonparametric model for $P_\star^0$ by $\mathcal{M}_\star$. Our objective is to obtain a lower bound on $\Psi(P_\star^{0,F})$ based on assumptions of how $P_\star^{0,F}$ is related to the distributions $P_1, \ldots, P_S$ from the completed trials. We refer to these assumptions as bridging assumptions. Because we only make assumptions that are biologically justifiable in a wide variety of examples, we are often unable to identify $\Psi(P_\star^{0,F})$ with any parameter mapping of the collection of distributions $\mathcal{P}^0 \triangleq (P_\star^0, P_1^0, \ldots, P_S^0)$. We are, however, able to identify a parameter mapping that lower bounds this quantity under our bridging assumptions. Obtaining such bounds is the objective of the partial identification literature (Manski, 2003). A consequence of our partial bridging is that our estimator of the lower bound does not generally converge to $\Psi(P_\star^{0,F})$ even as $n_\star, n_1, \ldots, n_S$ all grow to infinity.

Before proceeding, we introduce some notation. We let $\mathcal{P}^0 \triangleq (P_\star^0, P_1^0, \ldots, P_S^0)$ and $n \triangleq n_\star + \sum_{s=1}^S n_s$. We treat each of the sample sizes n_s as deterministic. We let $\mathcal{M} \triangleq \mathcal{M}_\star \times \prod_{s=1}^S \mathcal{M}_s$ denote the model for $\mathcal{P}^0$. For members P_s' of the general collection $\mathcal{P}' \triangleq \{P_\star', P_1', \ldots, P_S'\}$, we let $\mathbb{E}_s'$ denote the expectation

under $\mathcal{P}'$. We also let $\mathcal{S} = \{\star, 1, \ldots, S\}$. For Trials $s = 1, \ldots, S$, we define the conditional vaccine efficacy as follows

$$\mathrm{VE}_s^0(w) \triangleq 1 - \frac{\mathbb{E}_s^0[Y|A=1,w]}{\mathbb{E}_s^0[Y|A=0,w]}.$$

We define the unidentifiable conditional vaccine efficacy in population $\star$ similarly, with the expectations under P_s^0 above replaced by expectations under $P_\star^{0,F}$.

We now define a lower bound on the curve $w \mapsto \mathrm{VE}_\star^{0,F}(w)$. This lower bound is defined using data from the completed trials $s = 1, \ldots, S$, and also possibly vaccinated/unvaccinated conditional risks from a user-defined pseudo-population, which we denote $s = 0$. This pseudo-population can be used to make our vaccine efficacy curve assumption more plausible. Let $\mathtt{v}_s : \mathcal{W} \to [0,1]$, $s = 0, 1, \ldots, S$, be a collection of functions satisfying the convexity constraint $\sum_{s=0}^S \mathtt{v}_s(w) = 1$ for all w. Let $d : \mathcal{A} \times \mathcal{W} \to [0,1]$ be a function for which $d(0,w)$ and $d(1,w)$ respectively represent the unvaccinated and vaccinated risk in the user-defined pseudo-trial. Define

$$\mathrm{VE}_-^0(w) \triangleq 1 - \frac{\mathtt{v}_0(w) d(1,w) + \sum_{s=1}^S \mathtt{v}_s(w)\, \mathbb{E}_s[Y|A=1,w]}{\mathtt{v}_0(w) d(0,w) + \sum_{s=1}^S \mathtt{v}_s(w)\, \mathbb{E}_s[Y|A=0,w]}.$$

For each w, the above expression gives weight $\mathtt{v}_s(w)$ for the pseudo-trial $s = 0$ and each completed trial $s = 1, \ldots, S$. That is, $\mathtt{v}_s(w)$ indicates the hypothetical size of the stratum of w in trial s relative to the stratum of w in the other trials and the pseudo-trial. If $\mathtt{v}_0 \equiv 0$, then the pseudo-trial gets zero weight, whereas if $\mathtt{v}_0(w)$ is large for many w, then the vaccinated and unvaccinated risks $d(1,w)$ and $d(0,w)$ in the pseudo-trial play a major role in determining $w \mapsto \mathrm{VE}_-^0(w)$. Our first bridging assumption states that $w \mapsto \mathrm{VE}_-^0(w)$ lower bounds $w \mapsto \mathrm{VE}_\star^{0,F}(w)$.

B1) $P_\star^0\{\mathrm{VE}_\star^{0,F}(W) \geq \mathrm{VE}_-^0(W)\} = 1$.

If $\mathrm{VE}_\star^{0,F} \equiv \mathrm{VE}_1^0 \equiv \ldots \equiv \mathrm{VE}_S^0$, then the stronger result $P_\star^0\{\mathrm{VE}_\star^{0,F}(W) = \mathrm{VE}_-^0(W)\} = 1$ holds for any $w \mapsto \mathrm{VE}_-^0(w)$ for which the pseudo-trial is given zero weight, i.e. any $\mathtt{v}_0, \mathtt{v}_1, \ldots, \mathtt{v}_S$ satisfying $\mathtt{v}_0 \equiv 0$. We give a brief derivation of this result in Section 4.2.

While the above conditions are helpful for bridging the marginal vaccine efficacy from the completed trials to population $\star$, they do not allow one to uniquely identify this marginal vaccine efficacy, even when the conditional efficacy is point identifiable in the sense that $P_\star^0\{\mathrm{VE}_\star^{0,F}(W) = \mathrm{VE}_-^0(W)\} = 1$.

In particular, we still need to understand the behavior of the conditional risk among unvaccinated individuals in population $\star$. While getting a reasonable estimate of this baseline risk may be difficult in practice, experts may be able to give bounds on how small or large this risk may be. A flexible way of communicating this expert knowledge is as follows. Let $\mathsf{u}_0, \mathsf{u}_1, \ldots, \mathsf{u}_S$ be an expert-specified set of functions mapping from $\mathcal{W}$ to $\mathbb{R}$. Define the upper bound $\upsilon^0 \triangleq \Upsilon(\mathcal{P}^0)$ on the unvaccinated risk pointwise by

$$\upsilon^0(w) \triangleq \mathsf{u}_0(w) + \sum_{s=1}^{S} \mathsf{u}_s(w)\, \mathbb{E}_s^0[Y|A=0,w].$$

For the lower bound, let $\ell_0, \ell_1, \ldots, \ell_S$ be an expert-specified set of functions mapping from $\mathcal{W}$ to $\mathbb{R}$. The lower bound $\lambda^0 \triangleq \Lambda^{\mathcal{P}^0}$ on the unvaccinated risk is defined pointwise by

$$\lambda^0(w) \triangleq \ell_0(w) + \sum_{s=1}^{S} \ell_s(w)\, \mathbb{E}_s^0[Y|A=0,w],$$

where for each w we require that $\lambda^0(w) \leq \upsilon^0(w)$. We assume throughout that $\upsilon^0(w) - \lambda^0(w)$ is uniformly bounded below by some $\delta > 0$. Forcing $\delta > 0$ will not prove to be restrictive because Υ only needs to serve as an upper bound for the conditional baseline risk (known to fall in the closed unit interval). Having $\delta > 0$ ensures that point identification of the vaccine efficacy $\Psi(P_\star^{0,F})$ is impossible when the vaccine efficacy curve is nonconstant because we cannot identify the baseline risk in population $\star$, i.e. $w \mapsto \mathbb{E}_\star^{0,F}[Y|A=0,w]$.

The baseline risk assumption can be written as follows.

B2) $\lambda^0(W) \leq \mathbb{E}_\star^{0,F}[Y|A=0,W] \leq \upsilon^0(W)$ with $P_\star^0$ probability one.

Our next bridging assumption essentially states that the support of population $\star$ must be contained in the support of the completed trials, though, as we explain below, is slightly more general than this stated condition.

B3) For each $s = 1, \ldots, S$ and each event E on $\mathcal{W}$, $P_s^0(E) = 0$ implies that $\mathbb{E}_\star^0[\max\{\ell_s(W), \mathsf{u}_s(W), \mathsf{v}_s(W)\}\, \mathbb{1}_E] = 0$.

The above is weaker than assuming that the marginal distribution of W under P_s^0 dominates $P_\star^0$ for each $s = 1, \ldots, S$ since $\ell_s(W)$, $\mathsf{u}_s(W)$, and $\mathsf{v}_s(W)$ may be selected to be 0 for populations where there is no support.

We also assume that $w \mapsto \mathrm{VE}_-^0(w)$ satisfies a boundedness condition.

B4) $w \mapsto \mathrm{VE}_-^0(w)$ is uniformly bounded on the support of $\mathcal{W}$.

The above ensures that integrals involving $w \mapsto \mathrm{VE}_{-}^{0}(w)$ make sense – while it could be replaced by a moment condition, the boundedness assumptions simplifies our proofs and seems to give up little since it is rare that $\mathrm{VE}_{-}^{0}(w)$ can be arbitrarily close to $-\infty$.

2.2 Partial bridging formula

We now present lower bounds on the vaccine efficacy $\Psi(P_{\star}^{0,F})$ in population $\star$. We first present a loose lower bound that is attainable under only the bridging assumptions discussed thus far, and we then add a final bridging assumption that will often yield a tighter bound.

Given the bridging assumptions presented thus far, the tightest obtainable bound on $\Psi(P_{\star}^{0,F})$ is given by the solution to the following optimization problem:

$$\begin{aligned} &\text{Minimize} \quad \frac{\mathbb{E}_{\star}^{0}[g(W)]}{\mathbb{E}_{\star}^{0}[f(W)]} \text{ in } f, g : \mathcal{W} \to [0,1] \\ &\text{subject to} \quad \lambda^{0}(W) \leq f(W) \leq \upsilon^{0}(W),\ 1 - \frac{g(W)}{f(W)} \geq \mathrm{VE}_{-}^{0}(W), \end{aligned}$$

where we use the convention that $x/0 = 0$ if $x = 0$ and $x/0 = +\infty$ if $x > 0$. We now argue that the solution to this problem is undesirably loose in general, and that a further restriction will generally yield a tighter bound. For simplicity, we give the argument in the special case where one is unwilling to assume a lower bound on the unvaccinated risk so that $\lambda^{0} \equiv 0$. We then provide an alternative optimization problem that will typically give a tighter (larger) lower bound, present a closed-form solution to this new optimization problem, and finally we show why the above optimization problem is undesirably loose even in the case that λ^{0} is not the constant function zero. To ease discussion, the remainder of this paragraph supposes that there exists a $w_{-} \in \mathcal{W}$ such that $P_{\star}^{0}\{\mathrm{VE}_{-}^{0}(W) = \mathrm{VE}_{-}^{0}(w_{-})\} > 0$ and $P_{\star}^{0}\{\mathrm{VE}_{-}^{0}(W) < \mathrm{VE}_{-}^{0}(w_{-})\} = 0$, i.e. that $\mathrm{VE}_{-}^{0}(w_{-})$ is equal to the minimal value of $\mathrm{VE}_{-}^{0}(w)$ on $\mathcal{W}$. In this case one can quickly see that a $w \mapsto f(w)$ optimizing the above is positive at w_{-} and is zero everywhere else and the corresponding $w \mapsto g(w)$ is equal to $f(w_{-})\,\mathrm{VE}_{-}^{0}(w_{-})$ at w_{-} and 0 everywhere else, so that at this f, g we have that $1 - \frac{\mathbb{E}_{\star}^{0}[g(W)]}{\mathbb{E}_{\star}^{0}[f(W)]} = \mathrm{VE}_{-}^{0}(w_{-})$.

To overcome this problem, we will add an interpretable constraint to the optimization problem. For each $\mu > 0$ in the unit interval, we give a lower bound on $\Psi(P_{\star}^{0,F})$ that is valid if μ is equal to the marginal unvaccinated risk $\mathbb{E}_{\star}^{0}[\mathbb{E}_{\star}^{0,F}[Y|A = 0, W]]$. Each μ-specific lower bound can then be interpreted as

valid provided the already stated bridging assumptions hold and μ is equal to the the marginal unvaccinated risk. We derive our lower bound for $\Psi(P_\star^{0,F})$ by finding the worst-case conditional unvaccinated risk, namely by solving the following optimization problem:

$$\text{Minimize} \quad \frac{\mathbb{E}_\star^0[g(W)]}{\mathbb{E}_\star^0[f(W)]} \text{ in } f, g : \mathcal{W} \to [0,1]$$

$$\text{subject to} \quad \mathbb{E}_\star^0[f(W)] = \mu,\ \lambda^0(W) \leq f(W) \leq \upsilon^0(W),\ 1 - \frac{g(W)}{f(W)} \geq \mathrm{VE}_{_}^0(W).$$

The solution to our earlier problem in the special case where $\lambda^0(w)$ is always zero, namely $f(w_{_}) > 0$ and $f(w) = 0$ for $w \neq w_{_}$, will fail to satisfy the constraint $\mu = \mathbb{E}_\star^0[f(W)]$ once μ is large enough.

We now introduce notation to express a solution to our refined optimization problem. We define all of these quantities at $\mathcal{P}^0$, but the definitions at general $\mathcal{P}' \in \mathcal{M}$ are completely analogous. For each $\beta : \mathcal{W} \to [0,1]$ and $w \in \mathcal{W}$, let

$$\mathrm{UR}^\beta(\mathcal{P}^0)(w) \triangleq \lambda^0(w) + \left[\upsilon^0(w) - \lambda^0(w)\right]\beta(w) \triangleq \mathrm{UR}^{\beta,0}(w).$$

The set of all unvaccinated risk functions allowed by our bounds λ^0, υ^0 is equal to $\{\mathrm{UR}^{\beta,0} : \beta\}$. For each $\mathrm{UR}^{\beta,0}$, the marginal unvaccinated risk is given by

$$\Omega^\beta(\mathcal{P}^0) \triangleq \mathbb{E}_\star^0\left[\mathrm{UR}^{\beta,0}(W)\right] \triangleq \omega^{\beta,0}.$$

We also define

$$\Gamma^\beta(\mathcal{P}^0) \triangleq \mathbb{E}_\star^0\left[\mathrm{UR}^{\beta,0}\,\mathrm{VE}_{_}^0(W)\right] \triangleq \gamma^{\beta,0}.$$

Often $\omega^{\beta,0} - \gamma^{\beta,0}$ can be interpreted as a marginal vaccinated risk, though there is not in general any guarantee that this quantity is bounded in $[0,1]$. We discuss this subtlety further following Lemma 1.

We now define $\beta^0 : \mathcal{W} \to [0,1]$, which we will shortly show to be a β indexing the worst-case unvaccinated risk. The function β^0 assigns unvaccinated risk according to the upper bound υ^0 to as many covariate strata with small $\mathrm{VE}_{_}^0(w)$ as allowed by the marginal unvaccinated risk constraint μ. This threshold is defined by

$$\Theta(\mathcal{P}^0) \triangleq \sup\left\{\theta \in \mathbb{R} : \omega^{w \mapsto \mathbb{1}_{\{\mathrm{VE}_{_}^0(w) < \theta\}}, 0} \leq \mu\right\} \triangleq \theta^0,$$

where $\sup \emptyset = -\infty$ by convention. Let $\beta_\eta \triangleq w \mapsto \mathbb{1}_{\{\mathrm{VE}_{_}^0(w) < \theta^0\}} + \eta\, \mathbb{1}_{\{\mathrm{VE}_{_}^0(w) = \theta^0\}}$, and define η^0 to be the smallest element of the set $\mathrm{argmin}_{\eta \in [0,1]} \left(\omega^{\beta_\eta, 0} - \mu\right)^2$.

Let $\beta^0 \triangleq \beta_{\eta^0}$. One can show that, if $\mathbb{E}^0_\star[\lambda^0(W)] \leq \mu \leq \mathbb{E}^0_\star[\upsilon^0(W)]$, then $\omega^{\beta^0,0} = \mu$.

Finally, we define the bridging parameter that we will aim to estimate at $\mathcal{P}^0$, and note that the definition at general $\mathcal{P}'$ is entirely analogous:

$$\Phi(\mathcal{P}^0) \triangleq \frac{\Gamma^{\beta^0}(\mathcal{P}^0)}{\Omega^{\beta^0}(\mathcal{P}^0)} = \frac{\gamma^{\beta^0,0}}{\omega^{\beta^0,0}} \triangleq \phi^0,$$

where we remind the reader of the dependence of β^0 on $\mathcal{P}^0$. We now establish that ϕ^0 provides a valid lower bound for $\Psi(P^{0,F}_\star)$.

Lemma 1. *Suppose B1, B2, B3, and B4 hold. If* $\mu = \mathbb{E}^0_\star[\mathbb{E}^{0,F}_\star[Y|A=0,W]]$, *then* $\Psi(P^{0,F}_\star) \geq \phi^0$.

The proof of Lemma 1 is given in Appendix B.1. Typically the lower bound ϕ^0 will be attainable, in the sense that there exists a distribution for population $\star$ satisfying B1, B2, and $\mu = \mathbb{E}^0_\star[\mathbb{E}^{0,F}_\star[Y|A=0,W]]$ such that the marginal vaccine efficacy is equal to ϕ^0. In this case ϕ^0 is the solution to the refined optimization problem presented earlier in this section. In particular, it will often be the case that a distribution with marginal vaccine efficacy ϕ^0 is that with unvaccinated risk $\mathbb{E}^{0,F}_\star[Y|A=0,w] = \mathrm{UR}^{\beta^0,0}(w)$ and vaccine efficacy curve $w \mapsto \mathrm{VE}^0_-(w)$. The only time no such distribution exists is when the claimed vaccinated risk $w \mapsto \mathrm{UR}^{\beta^0,0}(w)[1-\mathrm{VE}^0_-(w)]$ fails to obey the bounds of the model, which may be larger than 1 if (i) $\mathrm{UR}^{\beta^0,0}(w)$ is large and $1-\mathrm{VE}^0_-(w) > 0$ or (ii) $1-\mathrm{VE}^0_-(w)$ is large. In practice it is unlikely that either of these quantities will be large: the unvaccinated risk will generally be small for rare outcomes, and the vaccine efficacy within a biomarker stratum will rarely be extremely negative, as this indicates that the vaccine is extremely harmful within this stratum and represents a situation where the bridging application would likely not be of interest and hence the method would not be applied.

Remark 1. Returning to our first optimization problem in this section, i.e. the optimization problem that did not include the constraint $\mathbb{E}^0_\star[f(W)] = \mu$, we see that the minimum occurs by minimizing the solution to the latter problem over all values of $\mu > 0$. We then see that the solution to the first optimization problem is equal to the lowest point on our μ-dependent curve, so that indeed the added restriction will generally improve our bound. For this observation to be true, one might anticipate that the reason that the solution to the first optimization problem is $\mathrm{VE}_-(w_-)$ in the case where the unvaccinated risk lower bound is the constant function zero, i.e. $\lambda^0 \equiv 0$, is that ϕ^0_μ is monotonic in μ,

where we write ϕ^0_μ to emphasize the dependence of ϕ^0 on μ. Indeed, we show in Lemma A.1 in Appendix B.1 that ϕ^0_μ is monotonically nondecreasing in μ in this case. This observation enriches the interpretation of the lower bound $\phi^0_{\mu'}$ when $\lambda^0 \equiv 0$, since in this case $\phi^0_{\mu'}$ can be interpreted as a valid lower bound for the marginal vaccine efficacy provided $\mu' \leq \mu = \mathbb{E}^0_\star[\mathbb{E}^{0,F}_\star[Y|A=0,W]]$ and the conditions of Lemma 1 hold at μ: as $\phi^0_{\mu'} \leq \phi^0_\mu$ and ϕ^0_μ lower bounds the vaccine efficacy, $\phi^0_{\mu'}$ must also lower bound the vaccine efficacy. Thus, in this special case, $\phi^0_{\mu'}$ is a valid lower bound whenever $\mu' \leq \mathbb{E}^0_\star[\mathbb{E}^{0,F}_\star[Y|A=0,W]]$. □

3 First-order expansion of lower bound

In this section, we present a result that we will use to derive a first-order expansion of the parameter Φ. This expansion plays a key role in our estimation procedure. Before presenting this result, we quickly define a gradient for a general parameter $\Pi : \mathcal{M} \to \mathbb{R}$ in the $S+1$ sample problem. We only define this and other gradients in this section at $\mathcal{P}^0$, but the extension to a general $\mathcal{P}'$ only requires notational changes. For each $s \in \mathcal{S}$, let $h_s : \mathcal{O}_s \to \mathbb{R}$ satisfy $\mathbb{E}_s[h_s(O_s)] = 0$ and $\sup_{o_s} |h_s(o_s)| \leq 1$. For $\epsilon \in (-1,1)$, define $\frac{dP^\epsilon_s}{dP^0_s}(o_s) = 1 + \epsilon h_s(o_s)$, $s \in \mathcal{S}$. Let $\mathcal{P}^\epsilon \triangleq (P^\epsilon_\star, P^\epsilon_1, \ldots, P^\epsilon_S)$. We call Π pathwise differentiable if, for $s \in \mathcal{S}$, there exist functions $\nabla\Pi^0_s \in L^2(P^0_s)$ satisfying

$$\frac{d}{d\epsilon}\Pi(\mathcal{P}^\epsilon)\bigg|_{\epsilon=0} = \sum_{s\in\mathcal{S}} \int \nabla\Pi^0_s(o_s) h_s(o_s) dP^0_s(o_s).$$

Our pathwise differentiability result will require one of the following three regularity conditions on the marginal distribution of W in population $\star$. The first is given below.

C1) $\mathbb{E}^0_\star[\lambda^0(W)] < \mu < \mathbb{E}^0_\star[\upsilon^0(W)]$ and, for all θ in a neighborhood of θ^0,

$$\begin{aligned} 0 &< \liminf_{t\to 0} \frac{P^0_\star\{\mathrm{VE}^0_-(W) < \theta + t\} - P^0_\star\{\mathrm{VE}^0_-(W) < \theta\}}{t} \\ &\leq \limsup_{t\to 0} \frac{P^0_\star\{\mathrm{VE}^0_-(W) < \theta + t\} - P^0_\star\{\mathrm{VE}^0_-(W) < \theta\}}{t} < \infty. \end{aligned}$$

The above implies both that θ^0 is finite and that $\omega^{\beta^0,0} = \mu$.

For the choice of μ to be feasible, i.e. for it to be possible that $P^{0,F}_\star \in \mathcal{M}_\star$ satisfies both B2 and $\mathbb{E}^0_\star[\mathbb{E}^{0,F}_\star[Y|A=0,W]] = \mu$, we generally need that

$\mathbb{E}^0_\star[\lambda^0(W)] \leq \mu \leq \mathbb{E}^0_\star[\upsilon^0(W)]$: one cannot otherwise have both $\mathbb{E}^0_\star[\mathbb{E}^{0,F}_\star[Y|A = 0, W]] = \mu$ and $\lambda^0(w) \leq \mathbb{E}^{0,F}_\star[Y|A = 0, w] \leq \upsilon^0(w)$ for all w. Nonetheless, it is useful to understand the first-order behavior of the parameter Φ for all values of μ since the marginal distribution of $W \sim P^0_\star$ must be estimated in practice. Thus, we offer conditions for both the case that μ is so large that it violates the upper bound on $\mathbb{E}^{0,F}_\star[Y|A = 0, w]$, i.e. $\mu > \mathbb{E}^0_\star[\upsilon^0(W)]$, and then for the case that μ is so small that it violates the lower bound on $\mathbb{E}^{0,F}_\star[Y|A = 0, w]$, i.e. $\mu < \mathbb{E}^0_\star[\lambda^0(W)]$.

C1') The upper bound is too small: $\mathbb{E}^0_\star[\upsilon^0(W)] < \mu$.

C1") The lower bound is too large: $\mathbb{E}^0_\star[\lambda^0(W)] > \mu$.

Remark 2. None of these three conditions (C1, C1', C1") allow $P^0_\star(\mathrm{VE}^0_-(W) = \theta^0) > 0$. This is closely related to the non-pathwise differentiability of many parameters of interest in personalized medicine under so-called exceptional laws (Robins, 2004; Luedtke and van der Laan, 2016a), i.e. distributions for which the conditional average treatment effect is zero in some positive probability stratum of measured covariates. The condition that $P^0_\star\{\mathrm{VE}_-(W) = \theta^0\} = 0$ may be unlikely to hold in settings where θ^0 is below the lower limit of quantification of W, namely because small values of W indicate a true value of zero for the biomarker and any deviation from zero is due to noise. If θ^0 falls below the lower limit if quantification, then we expect that $P^0_\star\{\mathrm{VE}^0_-(W) = \theta^0\} > 0$ and indeed we will have no guarantee that $\Psi(P^{0,F}_\star) \geq \phi^0$. In Appendix E.2, we describe an alternative bridging parameter that requires that W is univariate and VE^0_- is monotonic rather than that $P^0_\star\{\mathrm{VE}^0_-(W) = \theta^0\} > 0$.

Finally, we note that, while the conditions of Theorem 2 are sufficient for the pathwise differentiability of Φ, they are not necessary. For example, if W is discrete and takes on a finite number of values, then one can give conditions under which Φ is pathwise differentiable even if $P^0_\star(\mathrm{VE}^0_-(W) = \theta^0) > 0$, though these conditions still appear to require that only one $w \in \mathcal{W}$ satisfies $\mathrm{VE}^0_-(w) = \theta^0$. We do not consider this subtle case further in this work. □

The following objects, defined for each $s = 1, \ldots, S$, will be useful for establishing the gradient of Φ:

$$D^{\beta,\mathcal{P}^0}_{\mathrm{UR},s}(o_s) \triangleq [\ell_s(w) + \beta(w)\{\mathtt{u}_s(w) - \ell_s(w)\}] \frac{\mathbb{1}_{\{a=0\}}}{P^0_s(a|w)} \left(y - \mathbb{E}^0_s[Y|a, w]\right), \ \beta : \mathcal{W} \to \mathbb{R},$$

$$D^{\mathcal{P}^0}_{\mathrm{VE},s}(o_s) \triangleq \frac{\mathtt{v}_s(w)\left[\mathbb{1}_{\{a=0\}}\,\mathrm{VE}^0_-(w) - \mathbb{1}_{\{a=1\}}\right]}{P^0_s(a|w)\left[\mathtt{v}_0(w)d(0,w) + \sum_{s'=1}^S \mathtt{v}_{s'}(w)\,\mathbb{E}^0_{s'}[Y|A = 0, w]\right]}(y - \mathbb{E}^0_s[Y|a, w]). \tag{1}$$

For ease of notation, we let $D^{\beta,0}_{\mathrm{UR},s} \triangleq D^{\beta,\mathcal{P}^0}_{\mathrm{UR},s}$ and $D^{0}_{\mathrm{VE},s} \triangleq D^{\mathcal{P}^0}_{\mathrm{VE},s}$. For any $\beta : \mathcal{W} \to [0,1]$, define

$$
\begin{aligned}
\nabla\Omega^{\beta,0}_{\star}(o_\star) &= \mathrm{UR}^{\beta,0}(w) - \omega^{\beta,0}, \\
\nabla\Omega^{\beta,0}_{s}(o_s) &= \frac{dP^0_\star}{dP^0_s}(w) D^{\beta,0}_{\mathrm{UR},s}(o_s), && s = 1,\ldots,S, \\
\nabla\Gamma^{\beta,0}_{\star}(w) &\triangleq \mathrm{UR}^{\beta,0}(w)\,\mathrm{VE}^0_{-}(w) - \Gamma^{\beta}(\mathcal{P}^0), \\
\nabla\Gamma^{\beta,0}_{s}(o_s) &\triangleq \frac{dP^0_\star}{dP^0_s}(w)\left[D^{\beta,0}_{\mathrm{UR},s}(o_s)\,\mathrm{VE}^0_{-}(w) + \mathrm{UR}^{\beta,0}(w) D^{0}_{\mathrm{VE},s}(o_s)\right], \; s = 1,\ldots,S.
\end{aligned}
$$

We now give a theorem establishing that the parameter Φ is pathwise differentiable at $\mathcal{P}^0$.

Theorem 2. *If B3 and B4 hold and either C1, C1', or C1" holds, then Φ is pathwise differentiable and, for each $s \in \mathcal{S}$, the P^0_s gradient is given by*

$$
\nabla\Phi^0_s(o_s) = \begin{cases} \dfrac{\nabla\Gamma^{\beta^0,0}_s(o_s)}{\omega^{\beta^0,0}} - \theta^0 \dfrac{\nabla\Omega^{\beta^0,0}_s(o_s)}{\omega^{\beta^0,0}}, & \text{if } \mathbb{E}^0_\star[\lambda^0(W)] < \mu < \mathbb{E}^0_\star[\upsilon^0(W)], \\ \dfrac{\nabla\Gamma^{\beta^0,0}_s(o_s)}{\omega^{\beta^0,0}} - \gamma^{\beta^0,0} \dfrac{\nabla\Omega^{\beta^0,0}_s(o_s)}{[\omega^{\beta^0,0}]^2}, & \text{otherwise.} \end{cases}
$$

The proof of Theorem 2 is given in Appendix B.2.

Remark 3. Though we have assumed that $\upsilon^0(w) - \lambda^0(w)$ is bounded away from zero, we will now briefly remark on a violation of this assumption, namely the case that $\ell_s \equiv \mathsf{u}_s$ so that $\upsilon^0 \equiv \lambda^0$. In this case the gradients of Φ are given by the same expression as in Theorem 2 for the case where $\theta^0 = +\infty$ so that $P^0_\star(\mathrm{VE}^0_-(W) < \theta^0) = 1$. The estimation scheme that we present in the next section will remain valid even when $\upsilon^0 \equiv \lambda^0$ provided one uses the appropriate gradients, namely the gradients from Theorem 2 when $\theta^0 = +\infty$, when constructing the confidence lower bound. If $\mathrm{VE}^{0,F}_\star = \mathrm{VE}^0_-$, then the marginal vaccine efficacy in population $\star$ is point identified, and so our procedure is analogous to the earlier work by Rudolph and van der Laan (2016), with the slight distinction that we focus on a multiplicative rather than additive parameter. □

4 Estimation

4.1 Conditional vaccine efficacy not necessarily constant across $s = 1, \ldots, S$

We first consider the case where we do not assume that $\mathrm{VE}_1^0 \equiv \ldots \equiv \mathrm{VE}_S^0$, but rather make the weaker assumption B1 for a prespecified set of functions $\mathsf{v}_0, \mathsf{v}_1, \ldots, \mathsf{v}_S$.

4.1.1 Estimator. Below and throughout, we let Q_s^n denote the empirical distribution of the observations $O_s[1], \ldots, O_s[n_s]$ for each $s \in \mathcal{S}$. We also denote the expectation operator under $\mathcal{P}^n$ by $\mathbb{E}_s^n$. We also refer to parameters evaluated at the collection of distributions $\mathcal{P}^n$ rather than parameters evaluated at $\mathcal{P}^0$ by replacing the superscript zero by superscript n, e.g. we write $\upsilon^n \triangleq \Upsilon(\mathcal{P}^n)$ rather than $\upsilon^0 \triangleq \Upsilon(\mathcal{P}^0)$. We do the same for gradients, e.g. $\nabla\Phi_s^n$ rather than $\nabla\Phi_s^0$. When we define objects that may be confused with parameter mappings applied directly to $\mathcal{P}^n$, we add a hat: in particular, we will define $\widehat{\beta}^n$, $\widehat{\eta}^n$, $\widehat{\theta}^n$, $\widehat{\omega}^{\widehat{\beta}^n,n}$, $\widehat{\gamma}^{\widehat{\beta}^n,n}$, and $\widehat{\phi}^n$.

Below we present an estimation scheme to be used when either (i) the chosen $\ell_s(w)$ is a constant multiple of the chosen $\mathsf{u}_s(w)$, where the multiple is independent of s and w, or (ii) the chosen $\ell_s \equiv 0$ for all s so that $\lambda^0 \equiv 0$. In Appendix D, we present alternatives to the targeted minimum loss-based (TMLE) steps 2, 3, and 4 that replace the univariate logistic regression by a bivariate logistic regression. The bivariate regression is inappropriate when (i) holds because the two predictors in the proposed logistic regression will be perfectly correlated, and so, while analytically correct, the method may encounter numerical challenges in practice. When (ii) holds running the bivariate regression presented in the appendix is simply unnecessary because the latter covariate is always zero. The proof of asymptotic linearity of our estimation scheme assumes that the user runs the appropriate estimation scheme, either the below or that in Appendix D.

1. Let $(a, w) \mapsto \mathbb{E}_s^{n,\mathrm{init}}[Y|a, w]$, $w \mapsto P_s^{n,\mathrm{init}}(A = 1|w)$, and $w \mapsto \frac{dP_\star^{n,\mathrm{init}}}{dP_s^{n,\mathrm{init}}}(w)$ represent estimates of $(a, w) \mapsto \mathbb{E}_s^0[Y|a, w]$, $w \mapsto P_s^0(A = 1|w)$, and $w \mapsto \frac{dP_\star^0}{dP_s^0}(w)$, respectively.

2. Fit a univariate logistic regression with outcome $(y_s[i] : s = 1, \ldots, S;\ i = 1, \ldots, n_s)$, covariate $\left(\frac{\mathbb{1}_{\{a_s[i]=0\}}\,\mathsf{u}_s(w)}{n_s P_s^{n,\mathrm{init}}(A=0|w_s[i])}\frac{dP_\star^{n,\mathrm{init}}}{dP_s^{n,\mathrm{init}}}(w_s[i]) : s = 1, \ldots, S;\ i = 1, \ldots, n_s\right)$, and

fixed, subject-level intercept $\left(\text{logit}\left(\mathbb{E}_s^{n,\text{init}}[Y|A=0,w_s[i]]\right) : s=1,\ldots,S;\ i=1,\ldots,n_s\right)$. Denote the fitted coefficient in front of the covariate by ϵ_n.

3. For each $s=1,\ldots,S$, let $(a,w)\mapsto \mathbb{E}_s^{n,\epsilon_n}[Y|a,w]$ denote the function

$$(a,w)\mapsto \text{logit}^{-1}\left[\text{logit}\left(\mathbb{E}_s^{n,\text{init}}[Y|a,w]\right)+\epsilon_n \frac{\mathbb{1}_{\{a=0\}}\,\mathfrak{u}_s(w)}{n_s P_s^{n,\text{init}}(A=0|w)}\frac{dP_\star^{n,\text{init}}}{dP_s^{n,\text{init}}}(w)\right].$$

4. Let $\mathcal{P}^n=(P_\star^n,P_1^n,\ldots,P_s^n)$ denote any collection of distributions satisfying that, for all (a,w), $\mathbb{E}_s^n[Y|a,w]=\mathbb{E}_s^{n,\epsilon_n}[Y|a,w]$, $P_s^n(A=1|w)=P_s^{n,\text{init}}(A=1|w)$, and $\frac{dP_\star^n}{dP_s^n}(w)=\frac{dP_\star^{n,\text{init}}}{dP_s^{n,\text{init}}}(w)$.

5. For each $\beta:\mathcal{W}\to[0,1]$, let $\widehat{\omega}^{\beta,n}\triangleq\omega^{\beta,n}+\sum_{s\in\mathcal{S}}Q_s^n\nabla\Omega_s^{\beta,n}$, and note that $\widehat{\omega}^{\beta,n}$ rewrites as $Q_\star^n\,\text{UR}^{\beta,n}+\sum_{s=1}^S Q_s^n\nabla\Omega_s^{\beta,n}$.

6. Let $\widehat{\theta}^n\triangleq\sup\{\theta:\widehat{\omega}^{w\mapsto\mathbb{1}_{\{\text{VE}_-^n(w)<\theta\},n}}\leq\mu\}$, where $\sup\emptyset=-\infty$.

7. Let $\widehat{\eta}^n$ be the smallest element of the set $\text{argmin}_{\eta\in[0,1]}\left(\widehat{\omega}^{\beta_\eta,n}-\mu\right)^2$, where $\beta_\eta\triangleq w\mapsto\mathbb{1}_{\{\text{VE}_-^n(w)<\widehat{\theta}^n\}}+\eta\,\mathbb{1}_{\{\text{VE}_-^n(w)=\widehat{\theta}^n\}}$.

8. Let $\widehat{\beta}^n\triangleq\beta_{\widehat{\eta}^n}$.

9. Estimate $\gamma^{\widehat{\beta}^n,0}$ with

$$\begin{aligned}\widehat{\gamma}^{\widehat{\beta}^n,n}&\triangleq\gamma^{\widehat{\beta}^n,n}+\sum_{s\in\mathcal{S}}Q_s^n\nabla\Gamma_s^{\widehat{\beta}^n,n}\\&=n_\star^{-1}\sum_{i=1}^{n_\star}\text{UR}^{\widehat{\beta}^n,n}(w_\star[i])\,\text{VE}^n(w_\star[i])+\sum_{s=1}^S Q_s^n\nabla\Gamma_s^{\widehat{\beta}^n,n}.\end{aligned}$$

10. Estimate ϕ^0 with $\widehat{\phi}^n\triangleq\frac{\widehat{\gamma}^{\widehat{\beta}^n,n}}{\widehat{\omega}^{\widehat{\beta}^n,n}}$.

The initial estimates in Step 1 can be obtained by any methods deemed appropriate by the investigators. We encourage the use of data adaptive regression and machine learning techniques to estimate these features of $\mathcal{P}^0$, since the study of our estimator in the next section relies on consistency conditions given in Appendix A that are most plausible when one does not restrict the estimates in Step 1 to those obtained by classical parametric models.

Remark 4. The above represents a hybrid between a TMLE and a one-step estimator. When $\theta^0 = \pm\infty$, it uses a TMLE to estimate $\omega^{\beta^0,0}$ (Steps 2, 3, 4). It uses a one-step estimator for $\gamma^{\widehat{\beta}^n,0}$ (Step 9) and to correct the bias of $\omega^{\beta,n}$ for $\omega^{\beta,0}$ when defining $\widehat{\beta}^n$ (Step 5). For Step 9, one could specify a fluctuation of the initial estimate of $\mathcal{P}^0$ with score given by the sum of the empirical means of the $S+1$ gradients to replace the one-step estimator by a TMLE to have a full-fledged TMLE. One could analyze this estimator similarly to that of our hybrid estimator (van der Laan and Rose, 2011). The TMLE is preferred to a one-step estimator in many problems because it yields a plug-in estimator of the form $\gamma^{\widehat{\beta}^n,n}$ for some distribution $\mathcal{P}^n$, and is thus guaranteed to respect the known bounds on $\gamma^{\widehat{\beta}^n,0}$. Most fluctuation submodels indexing the TMLE appear to require iterating the fluctuation step until convergence. Alternatively, one could use the newly developed universal least favorable submodel, which has the advantage of always being fit in one step (van der Laan and Gruber, 2016). We leave the development of full-fledged (non-hybrid) TMLEs for this problem to future work. □

Remark 5. The upper and lower bounds υ^0 and λ^0 on the conditional unvaccinated risk imply bounds on the plausible range of values of μ. These bounds can be estimated using $\widehat{\omega}^{\beta,n}$, where β is taken to be the constant function returning one and zero, respectively. Under the same conditions as those given in the upcoming Theorem 3, one can show that, for each of these choices of β, $\widehat{\omega}^{\beta,n} - \omega^{\beta,0} \approx \sum_{s\in\mathcal{S}}(Q_s^n - P_s^0)\nabla\Omega_s^{\beta,0}$, so that confidence intervals for υ^0 and λ^0 can be developed using the analogous asymptotic normality results to those to be used to develop a confidence lower bound for ϕ^0 in Section 4.1.3. □

4.1.2 Asymptotic linearity of our estimator.

In the appendix we give several additional assumptions used to establish that the estimator $\widehat{\phi}^n$ satisfies a multiple sample version of asymptotic linearity, namely that

$$\widehat{\phi}^n - \phi^0 = \sum_{s\in\mathcal{S}}(Q_s^n - P_s^0)\nabla\Phi_s^0 + o_P(n_{\min}^{-1/2}), \tag{2}$$

where, for $n_{\min} \triangleq \min\{n_s : s \in \mathcal{S}\}$, we use $o_P(n_{\min}^{-1/2})$ to denote a random variable X_n that is a function of all of the n observations from Populations $\star, 1, \ldots, S$ with the property that, for each $\delta > 0$, the event $\{n_{\min}^{1/2}X_n > \delta\}$ has probability converging to zero whenever $n_{\min}$ grows to infinity. We discuss those conditions below the theorem.

The following theorem establishes that our estimator minus the truth behaves as a sum of empirical means across the $S+1$ data sets.

Theorem 3. *If B3, B4, C2, C3, C4, and C5 hold and either C1, C6, and C7 hold, C1' holds, or C1" holds, then* $\widehat{\phi}^n$ *satisfies (2).*

The proof of Theorem 3 is given in Appendix B.3. We now summarize the additional conditions used in the theorem. Condition C2 requires that the initial estimate of $\mathcal{P}^0$ is consistent and that certain remainder terms converge at a sufficient rate, namely $o_P(n_{\min}^{-1/2})$. Only requiring remainder terms to behave like $o_P(n_{\min}^{-1/2})$ makes our conditions on estimates of parameters from a full efficacy trial s more plausible if $n_s \gg n_{\min}$, since we typically expect these parameters to converge at an $n_s^{-1/2}$ rate in a parametric model and a slower rate in a nonparametric model, but the slower rate may still be much faster than $n_{\min}^{-1/2}$. Furthermore, the fact that W is univariate leads us to believe that, under some smoothness, we expect the nonparametric rate for estimating the needed functionals of $\mathcal{P}^n$ to be only slightly slower than the parametric rate. Conditions C3 and C4 represent an empirical process condition and mild consistency condition on the initial estimate $\mathcal{P}^n$. Condition C6 ensures that the estimate $\widehat{\beta}^n$ of β^0 satisfies $\widehat{\omega}^{\widehat{\beta}^n,n} \approx \mu$, where $\widehat{\omega}^{\widehat{\beta}^n,n}$ is a one-step estimate of $\omega^{\widehat{\beta}^n,0}$. This is to be expected under C1 given that $\omega^{\beta^0,0} = \mu$ in this case, much as we expect the empirical cumulative distribution function evaluated at the sample median to be equal to $1/2+o_P(n^{-1/2})$. In fact, we often expect $\widehat{\omega}^{\widehat{\beta}^n,n}$ to be exactly equal to μ, and this can be formally checked in any given application by looking at $\widehat{\omega}^{\widehat{\beta}^n,n}$. Condition C7 requires that VE_-^n estimates VE_-^0 sufficiently well so that $\mathrm{UR}^{\widehat{\beta}^n,0}$ is approximately equal to the worst-case unvaccinated risk function $\mathrm{UR}^{\beta^0,0}$. Appendix C relates this condition to the margin conditions appearing in the classification literature. Condition C5 simply states that the estimate of the vaccine efficacy curve $w \mapsto \mathrm{VE}_-^0(w)$ is bounded away from negative infinity, which is plausible given that $w \mapsto \mathrm{VE}_-^0(w)$ is bounded under B4.

4.1.3 Constructing a lower confidence bound. We now describe how Theorem 3 allows the construction of a lower confidence bound for the vaccine efficacy lower bound ϕ^0. For all $s \in \mathcal{S}$, let $\sigma_s^2 \triangleq \mathbb{E}_s^0[\nabla\Phi_s^0(O)^2]$, i.e. the variance of $\nabla\Phi_s^0(O)$ when $O \sim P_s^0$. Accepting slight notational overload, we define $\sigma_n^2 \triangleq \sum_{s\in\mathcal{S}} \frac{n_{\min}}{n_s}\sigma_s^2$. We propose using $\widehat{\phi}^n - 1.64 n_{\min}^{-1/2}\hat{\sigma}_n$ as a lower confidence bound for ϕ^0, where $\hat{\sigma}_n \triangleq \sum_{s\in\mathcal{S}} \frac{n_{\min}}{n_s}\hat{\sigma}_s^2$ for empirical estimates $\hat{\sigma}_s^2 \triangleq Q_s^n[\nabla\Phi_s^n - Q_s^n\nabla\Phi_s^n]^2$ of σ_s^2. Our lower confidence bound on ϕ^0 deviates from $\widehat{\phi}^n$ by the order $n_{\min}^{-1/2}\sigma_n$. As $\sigma_n < \sum_{s\in\mathcal{S}}\sigma_s^2 < \infty$ and $\widehat{\phi}^n$ deviates from ϕ^0 on the order $O(n_{\min}^{-1/2})$ by our asymptotic linearity result, the deviation of

our lower confidence bound from ϕ^0 is of the order of the inverse square root of the smallest trial. Given that the trials in Populations $1, \ldots, S$ will often have been efficacy studies, we will often expect the smallest sample to be of subjects drawn from population $\star$.

We can repeat the above procedure for a range of plausible μ values and report a graph of how the lower bound varies with the choice of μ. We remind the reader that the estimate $\widehat{\phi}^n$ and each $\hat{\sigma}_s^2$ rely on μ, and so these values must be recalculated for each choice of μ. Figure 2 from our simulation study suggests an informative way to display the vaccine efficacy lower bound, with the "average lower bound" y-axis (averaged across Monte Carlo replications of our simulation) replaced by the estimated lower bound. Given a range $\mathcal{U}$ of plausible values for μ, one can read a lower uncertainty interval (Vansteelandt et al., 2006) for the marginal vaccine efficacy $\Psi(P_\star^{0,F})$ in population $\star$ by looking at the smallest μ-specific confidence lower bound across all $\mu \in \mathcal{U}$.

We now justify our lower confidence bound. Note that (2) rewrites as

$$\frac{\sqrt{n_{\min}}[\widehat{\phi}^n - \phi^0]}{\sigma_n} = \sum_{s \in \mathcal{S}} \sum_{i=1}^{n_s} \frac{\sqrt{n_{\min}}}{n_s \sigma_n} \nabla \Phi_s^0(O_s[i]) + o_P(1). \tag{3}$$

For each s, observe that

$$\mathrm{Var}_s^0 \left[\frac{\sqrt{n_{\min}}}{n_s \sigma_n} \nabla \Phi_s(O_s[i]) \right] = \frac{1}{n_s} \frac{n_s^{-1} \sigma_s^2}{[n_\star^{-1} \sigma_\star^2 + \sum_{s'=1}^S n_{s'}^{-1} \sigma_{s'}^2]}.$$

The latter fraction on the right-hand side is bounded between 0 and 1. Thus the variance on the left shrinks at rate $O(n_s^{-1})$. The above and the independence of all of the observations also shows that the double sum on the right-hand side of (3) has variance 1.

These two facts readily allow one to show that the conditions of the Lindeberg central limit theorem hold, thereby establishing that $\frac{\sqrt{n_{\min}}[\widehat{\phi}^n - \phi^0]}{\sigma_n}$ converges to a standard normal distribution (Billingsley, 1999). It follows that a valid 95% level lower confidence bound for ϕ^0 is given by $\widehat{\phi}^n - 1.64 n_{\min}^{-1/2} \sigma_n$. In practice, σ_n is not known, but we can estimate it with $\hat{\sigma}_n$. Glivenko-Cantelli conditions on the functions $\nabla \Phi_s^n$ ensure that $\hat{\sigma}_n^2 \to \sigma_n^2$ in probability, and under this convergence Slutsky's theorem then ensures that a 95% lower bound for ϕ^0 is given by $\widehat{\phi}^n - 1.64 n_{\min}^{-1/2} \hat{\sigma}_n$.

4.1.4 Intuition for the estimator's construction and sketch of asymptotic linearity proof.

Our estimator is a hybrid of a TMLE and a one-step estimator. Before discussing our specific estimator, we give some intuition by

discussing one-step estimation of a general pathwise differentiable parameter $\Pi : \mathcal{M} \to \mathbb{R}$. Suppose that $\mathcal{P}^n$ is an initial estimate of $\mathcal{P}^0$ and we wish to estimate $\pi^0 \triangleq \Pi(\mathcal{P}^0)$. One possible estimate is obtained via plug-in estimation: $\pi^n \triangleq \Pi(\mathcal{P}^n)$. If $\mathcal{P}^n$ is a good estimate of $\mathcal{P}^0$, often in the sense that certain regression functions under $\mathcal{P}^n$ are close to the corresponding regressions under $\mathcal{P}^0$, then we generally expect that $\pi^n - \pi^0 \approx -\sum_{s\in\mathcal{S}} P_s^0 \nabla\Pi_s^n$, where $\nabla\Pi_s^n$ are gradients of Π evaluated at $\mathcal{P}^n$. It will often be the case that the right-hand side of this approximation is not $O_P(n_{\min}^{-1/2})$, and thus that π^n converges to π^0 at a slower than $n_{\min}^{-1/2}$ rate. In an effort to improve the initial estimate π^n, a one-step estimator adds the empirical version of $\sum_{s\in\mathcal{S}} P_s^0 \nabla\Pi_s^n$ to the initial estimate π^n, i.e. one uses as estimate $\widehat{\pi}^n \triangleq \pi^n + \sum_{s\in\mathcal{S}} Q_s^n \nabla\Pi_s^n$, so that $\widehat{\pi}^n - \pi^0 \approx \sum_{s\in\mathcal{S}} (Q_s^n - P_s^0)\nabla\Pi_s^n$. Under empirical process conditions, one can replace the gradients $\nabla\Pi_s^n$ under $\mathcal{P}^n$ by the gradients $\nabla\Pi_s^0$ under $\mathcal{P}^0$ so that $\widehat{\pi}^n - \pi^0 \approx \sum_{s\in\mathcal{S}} (Q_s^n - P_s^0)\nabla\Pi_s^0$ and it is possible to apply a central limit theorem to understand the behavior of a scaled version of the one-step estimator minus the truth. Typically TMLEs follow a similar derivation, but the initial estimate $\mathcal{P}^n$ of $\mathcal{P}^0$ is carefully selected so that the bias correction $\sum_{s\in\mathcal{S}} Q_s^n \nabla\Pi_s^0$ is small or zero.

The remainder of the discussion is focused on our proposed estimator. Steps 2, 3, and 4 are designed to ensure that the score $\sum_{s=1}^S Q_s^n \nabla\Omega_s^{\beta,n}$ is equal to zero when $\beta \equiv 1$, which is an important score equation to solve if $\widehat{\theta}^n = +\infty$. If each $\ell_s(w)$ is a constant multiple of $\mathfrak{u}_s(w)$ or if each $\ell_s \equiv 0$, then the same score equation is equal to zero when $\beta \equiv 0$, which is an important score equation to solve when $\widehat{\theta}^n = -\infty$. Otherwise, Steps 2', 3', and 4' in Appendix D ensure that the score equation is solved at both $\beta \equiv 0$ and $\beta \equiv 1$. If $\widehat{\theta} = +\infty$, then $\mathbb{1}_{\{\mathrm{VE}_-^n(w_\star[i]) < \widehat{\theta}^n\}} = 1$ for all $i = 1, \ldots, n_\star$ so that the preceding argument implies that $\sum_{s=1}^S Q_s^n \nabla\Omega_s^{\widehat{\beta}^n,n} = 0$ and Step 5 trivially yields that $Q_\star^n \nabla\Omega_\star^{\widehat{\beta}^n,n} = 0$. Similarly, $\widehat{\theta}^n = -\infty$ implies that $\sum_{s=1}^S Q_s^n \nabla\Omega_s^{\widehat{\beta}^n,n} = 0$.

We now discuss the choice of $\widehat{\beta}^n$. For each β, $\widehat{\omega}^{\beta,n}$ is a one-step estimator for $\omega^{\beta,0}$. It is therefore not surprising that, under the conditions given in Appendix A, $\widehat{\omega}^{\beta,n} = \omega^{\beta,0} + O_P(n_{\min}^{-1/2})$. If θ^0 is finite, then one can show $\widehat{\theta}^n \approx \theta^0$ (see Lemma A.14 in Appendix C). Furthermore, if μ is too large so that $\theta^0 = +\infty$, then our conditions imply that $\widehat{\theta}^n = \infty$ with probability approaching one (see Lemma A.9 in the Appendix B.3), and an analogous result holds if μ is too small.

Given an estimate of $\widehat{\beta}^n$, it remains to estimate ϕ^0. We show in Ap-

pendix B.3 that

$$\widehat{\phi}^n - \phi^0 \approx \frac{\widehat{\gamma}^{\widehat{\beta}^n,n} - \gamma^{\widehat{\beta}^n,0}}{\omega^{\beta^0,0}} - \frac{\gamma^{\beta^0,0}}{[\omega^{\beta^0,0}]^2}\left[\widehat{\omega}^{\widehat{\beta}^n,n} - \omega^{\beta^0,0}\right] + \frac{\gamma^{\widehat{\beta}^n,0} - \gamma^{\beta^0,0}}{\omega^{\beta^0,0}}.$$

For the first term on the right, we note that, $\widehat{\gamma}^{\beta,n}$ is a one-step estimator for $\gamma^{\beta,0}$ for each β, and so it is not surprising that this estimator is asymptotically linear for each β. Furthermore, the fact that $\widehat{\beta}^n$ converges to the fixed quantity β^0 suggests that these same asymptotic linearity statements should hold with β replaced by $\widehat{\beta}^n$ (can be formally shown via empirical process arguments). Under C1 and C6, $\widehat{\omega}^{\widehat{\beta}^n,n} \approx \mu = \omega^{\beta^0,0}$ so that the second term above is negligible. Under C1' or C1", the $\widehat{\beta}^n$ is correctly specified with probability approaching one so that analyzing the second term is like analyzing the one-step estimator $\widehat{\omega}^{\beta^0,n}$ of $\omega^{\beta^0,0}$, which is asymptotically linear under reasonable conditions. Additionally, C1' or C1" implies that the final term above is equal to zero with probability approaching one. For the final term above under C1, we note that

$$\gamma^{\widehat{\beta}^n,0} - \gamma^{\beta^0,0} = \beta^0\left[\omega^{\widehat{\beta}^n,0} - \omega^{\beta^0,0}\right] + \left\{\gamma^{\widehat{\beta}^n,0} - \gamma^{\beta^0,0} - \beta^0\left[\omega^{\widehat{\beta}^n,0} - \omega^{\beta^0,0}\right]\right\},$$

where the latter term is negligible by assumption C7 (though a more interpretable sufficient condition is given in Appendix C). We establish the behavior of $\omega^{\widehat{\beta}^n,0} - \omega^{\beta^0,0}$ by using that $\widehat{\omega}^{\widehat{\beta},n} \approx \mu = \omega^{\beta^0,0}$, and so it suffices to study the behavior of $-[\widehat{\omega}^{\widehat{\beta}^n,0} - \omega^{\widehat{\beta}^n,0}]$. The study of this term is standard since $\widehat{\omega}^{\beta,0}$ is a one-step estimator of the β-specific parameter $\omega^{\beta,0}$.

4.2 Conditional vaccine efficacy constant across $s = 1, \ldots, S$

The lower bound in the above section relies on the validity of the bridging assumption B1, which is indexed by d and $\mathtt{v}_0, \mathtt{v}_1, \ldots, \mathtt{v}_S$. Suppose that the following condition holds:

B5) The vaccine efficacy curve $w \mapsto \mathrm{VE}_s(w)$ is constant across trials $s = 1, \ldots, S$ and the chosen $\mathtt{v}_0 \equiv 0$.

Choosing $\mathtt{v}_0 \equiv 0$ indicates that the user believes that the vaccine efficacy curves in the completed trials lower bound the vaccine efficacy curve in population $\star$.

We now show that the above condition allows one to choose $\mathtt{v} \triangleq (\mathtt{v}_s : s = 1, \ldots, S)$ to maximize statistical efficiency for estimating our lower bound. In what follows we make the dependence of $w \mapsto \mathrm{VE}^0_-(w)$ on $\mathtt{v}$ explicit by writing

$\mathrm{VE}_{-}^{\mathtt{v},0}$. The result relies on the fact that $\mathrm{VE}_{-}^{\mathtt{v},0}$ is invariant in $\mathtt{v}$. In particular, for each w we have that

$$\begin{aligned}\mathrm{VE}_{-}^{\mathtt{v},0}(w) &= \frac{\sum_{s=1}^{S} \mathtt{v}_s(w)\, \mathbb{E}_s[Y|A=0,w]\, \mathrm{VE}_s^0(w)}{\sum_{s=1}^{S} \mathtt{v}_s(w)\, \mathbb{E}_s[Y|A=0,w]} \\ &= \mathrm{VE}_1^0(w) \frac{\sum_{s=1}^{S} \mathtt{v}_s(w)\, \mathbb{E}_s[Y|A=0,w]}{\sum_{s=1}^{S} \mathtt{v}_s(w)\, \mathbb{E}_s[Y|A=0,w]},\end{aligned}$$

which is equal to $\mathrm{VE}_1^0(w)$. Thus, under B5, if B1 holds for some set of functions $\mathtt{v}$, then it holds for all $\mathtt{v}$. While the above derivation implies that the lower bound parameter Φ is also invariant to $\mathtt{v}$, the same is not true of the gradients presented in Theorem 2. As the gradient determines the efficiency of our estimator, and therefore the width of our confidence interval, it follows that one can choose $\mathtt{v}$ to (approximately) maximize the efficiency of our procedure provided B5 holds. Let $\hat{\sigma}_n^2(v)$ and $\sigma_n^2(v)$ respectively denote the values of $\hat{\sigma}_n^2$ and σ_n^2 defined in Section 4.1.3, but now making the dependence on $\mathtt{v}$ explicit in the notation. If the consistency conditions of Theorem 3 hold uniformly over all $\mathtt{v}$ in some class $\mathcal{V}$, then one expects that

$$\sup_{v \in \mathcal{V}} |\hat{\sigma}_n^2(v) - \sigma_n^2(v)| \to 0. \tag{4}$$

As $\hat{\sigma}_n^2(v)$ can be written as an empirical mean of random functions, the primary condition that one needs to add to Theorem 3 for this convergence to be valid is that $\mathcal{V}$ is not too large, namely that functions in the sets in $\mathcal{V}$ belong to a Glivenko-Cantelli class. We will return to this requirement shortly. For now, suppose that we have selected a class $\mathcal{V}$ small enough so that (4) is plausible. One can then select $\mathtt{v}_n$ as the minimizer of $\hat{\sigma}_n^2(v)$ over $v \in \mathcal{V}$. If $\mathcal{V}$ satisfies a more restrictive condition, namely that the functions in the sets it contains are Donsker, then we expect that we can run the estimator described in Section 4.1.1 at the selected $\mathtt{v}_n$ and report the lower confidence bound as defined in Section 4.1.3. A simple suggestion for a class $\mathcal{V}$ that satisfies the conditions of the previous theorem is that $\mathcal{V}$ consists of constant functions, namely each $\mathtt{v}_s(w) \in [0,1]$ does not depend on w. One could alternatively parameterize $\mathtt{v}_s$ using, e.g., a linear logistic regression formulation so that $\mathtt{v}_s(w) \propto \mathrm{logit}^{-1}(\beta_0 + \beta_1 w)$ to further improve efficiency.

The arguments that we have provided here are only a sketch: more work would be needed to make them precise, though these arguments are fairly standard when studying M-estimators (van der Vaart and Wellner, 1996), and a detailed study specific to an estimator with a nuisance parameter selected to minimize the variance were given in Rubin and van der Laan (2008). A careful

analysis would show that, when B5 holds and $\mathcal{V}$ is not too large, the conditions needed to ensure the validity of the procedure discussed in this section are not much stronger than the conditions needed to establish the validity of a procedure at fixed $\mathtt{v}$. Indeed, in Section 5 we show via simulation that this procedure yields less conservative lower bounds while still maintaining nominal coverage.

5 Simulation

5.1 Simulation settings

We evaluated the performance of our method via simulation in `R` (R Core Team, 2014). We have $S = 2$ completed efficacy trials in our simulation. We first run our simulation without missingness in the biomarker W, and then we simulate data from a two-phase sampling scheme.

Let Z be a standard normal random variable. The marginal distribution of W has the same distribution as $\text{logit}^{-1}(Z-2)$ when $s = 1$, $\text{logit}^{-1}(5[Z-1])$ when $s = 2$, and $\text{logit}^{-1}(2[Z-1/2])$ when $s = \star$. All trials assign treatment with probability $1/2$, regardless of the value of baseline covariates. The vaccine efficacy is the same across Populations 1 and 2. In particular, both have vaccine efficacy

$$\text{VE}(w) = 1 - \frac{\text{logit}^{-1}\left(-1 - w - 3[0.3 + (w - 0.2)^+]\right)}{\text{logit}^{-1}\left(-1 - w\right)},$$

where x^+ denotes the positive part of x. As $\mathbb{E}_s^0[Y|A = 1, w]$ equals $[1 - \text{VE}(w)]\,\mathbb{E}_s^0[Y|A = 0, w]$, it suffices to define $\mathbb{E}_s^0[Y|A = 0, w]$ for each of the three trials. We let $\mathbb{E}_1^0[Y|A = 0, w] = \text{logit}^{-1}\left(-1 - w\right)$ and $\mathbb{E}_2^0[Y|A = 0, w] = \text{logit}^{-1}\left(-1\right) \approx 0.27$.

We consider two sample size settings, namely $(n_\star, n_1, n_2)$ equal to $(100, 2000, 2000)$ and $(200, 4000, 4000)$. We respectively refer to these sample sizes as the "Smaller Sample Size" and "Larger Sample Size" in our figures. We use three choices for the coefficients indexing our unvaccinated risk lower and upper bounds. The first setting, labeled "Loosest" in our figures, sets $\ell_s \equiv \mathtt{u}_s \equiv 0$, $s = 1, 2$, and $\ell_0 \equiv 0$, $\mathtt{u}_0 \equiv 1$. The second, labeled "Moderate" in our figures, sets $\ell_s \equiv 0.25$ and $\mathtt{u}_s \equiv 0.75$, $s = 1, 2$, and $\ell_0 \equiv \mathtt{u}_0 \equiv 0$. The third, labeled "Tight" in our figures, sets $\ell_s(w) = 0.4$ and $\mathtt{u}_s(w) = 0.6$, $s = 1, 2$, and $\ell_0 \equiv \mathtt{u}_0 \equiv 0$.

We suppose that the user *a priori* believes that $\text{VE}_\star^{0,F}(w) \geq \text{VE}_1(w)$ and that $\text{VE}_s(w)$ is invariant in s. In this case $\mathtt{v}_0(w) = 0$, while $w \mapsto \text{VE}_-^0(w)$ is in-

variant in the choice of $\mathtt{v}$. We implement the version of our estimator presented in the main text since $\ell_s(w)$ is a multiple of $\mathtt{u}_s(w)$ for all simulation settings, both for fixed $\mathtt{v}_1(w) = \mathtt{v}_2(w) = 1/2$, all w, indexing $w \mapsto \mathrm{VE}_{-}^{0}(w)$ and also for the procedure described in Section 4.2 that adaptively selects a w-invariant convex combination $\mathtt{v}_1(w), \mathtt{v}_2(2)$ to maximize the estimator's precision. We respectively refer to these estimation schemes as "Fixed" and "Adaptive" in our figures. The marginal distributions of the biomarker within each population are estimated using a ratio of kernel density estimates as provided by the `ks` package (Duong, 2016), with the option `unit.interval` set to `TRUE` so that the estimator respects the bounds on the biomarker W and the bandwidth selected according to the univariate plug-in selector presented in `?`. The ratio is then standardized so that the empirical mean of each s-specific estimate of $\frac{dP_\star^0}{dP_s^0}$ has empirical mean 1 within the sample of observations from Trial s. Although the probability of treatment given covariates is known in each trial, we estimate these quantities because of the known efficiency gains resulting from doing so (van der Laan and Robins, 2003). In particular, we regress A against W using the ensemble algorithm found in the `SuperLearner` package (van der Laan et al., 2007; Polley and van der Laan, 2013), and provide the algorithm with a candidate library containing `SL.glm` and `SL.glm.interaction`. We estimate the outcome regressions by regressing Y against A and W using the `SuperLearner` package, using a candidate library containing `SL.mean`, `SL.glm`, `SL.glm.interaction`, `SL.step.interaction`, `SL.gam`, and `SL.nnet`.

We then repeat our simulation with missingness in the biomarker in the completed efficacy trials via a 1:1 nested case:control sampling design. In particular, we suppose that biomarker values are observed on all m cases (participants with $Y = 1$), and, within each trial s, on a random sample of controls of size m. We run the method described in Appendix E, first using the known observation weights, and then estimating these weights by using the known unity weights among cases and, among controls, running a logistic regression of observation status against a linear interaction model of observation status against an indicator for belonging to Trial $s = 1, 2$ and an observed covariate $X = \mathrm{logit}^{-1}[\mathrm{logit}(W) + Z]$, where Z is a standard normal. The covariate X is meant to represent a biomarker that is inexpensive to measure on all participants, but that is highly predictive of the biomarker W needed for the bridging exercise.

The function VE_{-}^{0} used in our simulation is monotonic, and therefore the results in Appendix E.2 may be used to weaken the assumptions for our method's applicability. The results were nearly identical for this data generating distribution and so are omitted.

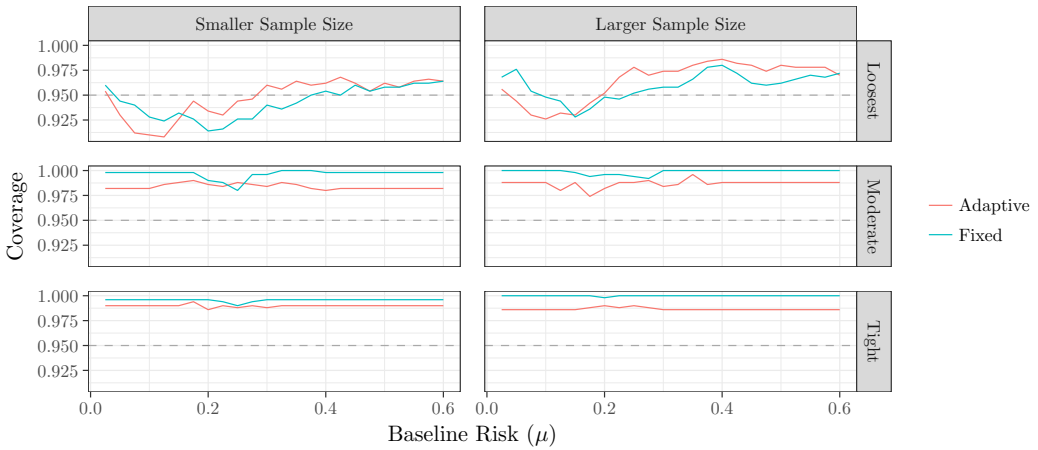

Figure 1: Coverage of our lower confidence bound for ϕ^0, i.e. the lower bound on the vaccine efficacy, in our simulation. Conducted at both smaller and larger sample sizes, respectively with $(n_\star, n_1, n_2)$ equal to $(100, 2000, 2000)$ and $(200, 4000, 4000)$, and for different choices of ℓ_s and u_s, determining the tightness of the unvaccinated risk bounds. Horizontal dashed lines drawn at 95% coverage.

5.2 Simulation results

Figure 1 demonstrates that our lower confidence bounds achieve approximately 95% or better coverage for the lower bound on the vaccine efficacy across simulation settings and values of μ. The bound is generally conservative for the Moderate and Tight settings. For a given data generating distribution, there is one true value of the marginal unvaccinated risk μ_0, and the interpretation of our curve as a 95% lower confidence bound is valid if the lower bound coverage is approximately 95% when $\mu = \mu_0$. The methods that choose the convex combination in the $w \mapsto \mathrm{VE}^0_-(w)$ bridging parameter to minimize the standard error had comparable coverage than the methods that used a fixed convex combination.

Figure 2 demonstrates that the adaptive methods tend to have a slightly higher average lower bound than the fixed methods that use a convex combination of $(1/2, 1/2)$ for the VE bridging assumption. Given the approximately appropriate coverage of both methods for the (typically conservative) lower bound on the true vaccine efficacy, it seems prudent in practice to use the less conservative adaptive approach.

In Appendix F, we display analogous plots for the fixed two-phase sampling method. Results were similar for the adaptive method. These plots show that

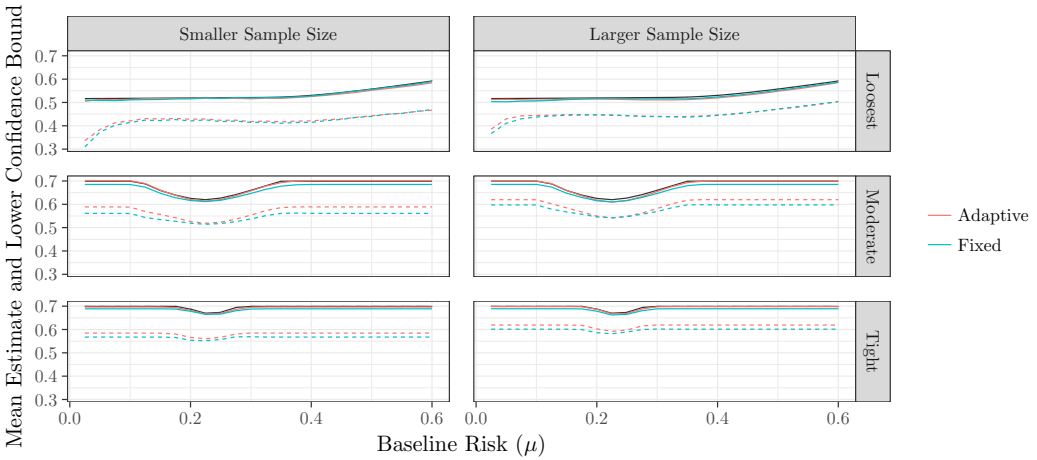

Figure 2: Average estimates (solid lines) and lower confidence bounds (dashed lines) for ϕ^0, i.e. the lower bound on the vaccine efficacy, in our simulation. Conducted at both smaller and larger sample sizes, respectively with $(n_\star, n_1, n_2)$ equal to $(100, 2000, 2000)$ and $(200, 4000, 4000)$, and for different choices of ℓ_s and u_s, determining the tightness of the unvaccinated risk bounds. Black trend lines denote true $(\ell_s, \mathsf{u}_s, \mu)$-specific lower bound.

the two-phase sampling procedure still maintains the desired coverage level, with a moderate loss in precision in our setting. This is not surprising given that the user has less data for the two-phase methods.

6 Discussion

We have presented a method for using data from completed efficacy trials to partially bridge the vaccine efficacy to a new population. We first developed conditions that identify a population level parameter representing a lower bound on the vaccine efficacy in the new population. We then provided a nonparametric estimator of this quantity that respects the fact that this is a multiple sample problem and enables the use of modern data adaptive regression and density estimation techniques to estimate the underlying baseline biomarker distribution, worst-case unvaccinated risk, and vaccine efficacy lower bound.

The validity of our population level lower bound on the vaccine efficacy in the new population relies on three main conditions. One of these conditions essentially states that the support of the key biomarker in the new setting is contained in that of the old efficacy trials so that it is possible to learn about

the new setting from these trials. Another states that the results from the earlier trials can be used to derive a lower bound of the conditional vaccine efficacy curve in the new setting. A special case of this condition is that the conditional vaccine efficacy curve is constant across trials. While these two conditions alone allow one to get a lower bound on the marginal vaccine efficacy, this lower bound can be very loose when the biomarker is truly a modifier of vaccine efficacy because it involves finding the worst case distribution for the conditional unvaccinated risk among the class of all such possible conditional risk distributions. We thus add a third condition, namely that the marginal unvaccinated risk is equal to a user-specified constant μ. Adding this constraint can greatly tighten the lower bound. The user can then report lower confidence bounds resulting from our procedure at a range of plausible values of μ.

In this work, we have assumed that one is able to define a lower bound on the vaccine efficacy conditional only on baseline covariates and not on (counterfactual) post-vaccination immune responses. There are advantages and disadvantages to this problem setup. The advantage is that, if the lower bound assumption is valid, investigators can partially bridge the vaccine efficacy to a new setting without vaccinating any individuals in the new population and subsequently waiting to measure the immune response. Avoiding running a phase I immunogenicity trial can save significant financial resources and can also accelerate the introduction of a potentially life-saving vaccine into the new population. The disadvantage of bridging using only using baseline covariates is that the post-vaccination immune response is often more strongly associated with vaccine efficacy, and so a partial bridging assumption that accounts for these responses may be more plausible or less conservative. Gilbert and Huang (2016) consider bridging within a principal stratification framework, thereby allowing for conditioning on counterfactual immune responses. In a future work, we will present a doubly robust partial bridging procedure that allows for conditioning on post-vaccination biomarkers.

Our vaccine efficacy lower bound depends on the marginal unvaccinated risk level μ. While in many settings subject-matter experts may be able to suggest a plausible range of values for μ, in other cases there may not be sufficient knowledge of the marginal risk for this to be possible. To deal with this issue, one could follow the suggestion of Gilbert and Huang (2016) to formally incorporate epidemiologic surveillance data to learn such a range. While there may not be biomarker and baseline risk information jointly available, one may at times have reasonably accurate active surveillance data available to estimate the marginal unvaccinated risk in the new population. In this case, one could estimate our μ parameter rather than report a range of plausible values.

One could then either treat this estimate as the truth, or modify our inferential procedure to incorporate the added uncertainty. While active surveillance data should be incorporated when it exists, often only data from national surveillance systems is available. These data may not be sufficiently accurate to inform plausible values of μ. For example, recent studies of national surveillance of dengue incidence have shown up to 19-fold underreporting compared to active surveillance (Nealon et al., 2016; Sarti et al., 2016).

The results in this paper readily extend to a large class of contrasts $f(\mu_1, \mu_0)$ between the marginal vaccinated risk μ_1 and unvaccinated risk μ_0. We have studied the multiplicative efficacy $f(\mu_1, \mu_0) = 1 - \mu_1/\mu_0$. If one is willing to make the three bridging assumptions made in this paper, namely a support condition, lower bounded conditional *multiplicative* vaccine efficacy, and specified unvaccinated risk, then obtaining a lower bound for more general $f(\mu_1, \mu_0)$ is straightforward. Suppose that, for each μ_0, $f(\mu_1, \mu_0)$ is sufficiently smooth and monotonically decreasing in μ_1 (the higher the risk among vaccinated individuals, the lower the efficacy). For a user-defined μ, one can then conservatively estimate $f(\mu_1, \mu)$ by upper bounding μ_1, which is indeed what we do in Section 2. When $\mu = \mu_0$, one then gets a lower bound for the efficacy parameter of interest. All typical contrasts, including additive and odds-ratio type contrasts, appear to satisfy this monotonicity condition.

Finally, we note that, while we have focused our discussion on the intervention A being a vaccination indicator, the results given in this paper apply immediately to any binary intervention A and multiplicative efficacy parameter.

Acknowledgements

This work was partially supported by the National Institute of Allergy and Infectious Disease at the National Institutes of Health under award numbers R37AI054165 and UM1 AI068635. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. The authors thank Ying Huang, Kara Rudolph, Ying Chen, Holly Janes, Dobromir Dimitrov, Jonathan Sugimoto, Laura Matrajt and Paul Edlefsen for their thoughtful comments on the early stages of this project.

References

J Y Audibert and A B Tsybakov. Fast learning rates for plug-in classifiers. *The Annals of Statistics*, 35(2):608–633, 2007.

E Bareinboim and J Pearl. Transportability of Causal Effects: Completeness Results. In *Twenty-Sixth AAAI Conference on Artificial Intelligence*, 2012.

E Bareinboim and J Pearl. A general algorithm for deciding transportability of experimental results. *Journal of causal Inference*, 1(1):107–134, 2013.

E Bareinboim and J Pearl. Transportability from multiple environments with limited experiments: Completeness results. In *Advances in Neural Information Processing Systems*, pages 280–288, 2014.

P Billingsley. *Convergence of probability measures*. Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons Inc., New York, second edition, 1999. ISBN 0-471-19745-9.

N E Breslow. Statistics in Epidemiology: the Case-Control Study. *J Am Stat Soc*, 91:14–28, 1996.

M R Capeding, N H Tran, S R S Hadinegoro, H I H J M Ismail, T Chotpitayasunondh, M N Chua, C Q Luong, K Rusmil, D N Wirawan, R Nallusamy, and Others. Clinical efficacy and safety of a novel tetravalent dengue vaccine in healthy children in Asia: a phase 3, randomised, observer-masked, placebo-controlled trial. *The Lancet*, 384(9951):1358–1365, 2014.

S R Cole and E A Stuart. Generalizing evidence from randomized clinical trials to target populations the actg 320 trial. *American journal of epidemiology*, 172(1):107–115, 2010.

T Duong. *ks: Kernel Smoothing*, 2016. URL https://cran.r-project.org/package=ks.

T R Fleming and J H Powers. Biomarkers and surrogate endpoints in clinical trials. *Statistics in medicine*, 31(25):2973–2984, 2012.

P B Gilbert and Y Huang. Predicting Overall Vaccine Efficacy in a New Setting by Re-calibrating Baseline Covariate and Intermediate Response Endpoint Effect Modifiers of Type-Specific Vaccine Efficacy. *Epidemiologic Methods*, 2016.

P B Gilbert and M G Hudgens. Evaluating candidate principal surrogate endpoints. *Biometrics*, 64(4):1146–1154, 2008.

P B Gilbert, M G Hudgens, and J Wolfson. Commentary on" Principal Stratification–a Goal or a Tool?" by Judea Pearl. *The international journal of biostatistics*, 7(1):1–15, 2011.

J L Horowitz and C F Manski. Nonparametric analysis of randomized experiments with missing covariate and outcome data. *Journal of the American statistical Association*, 95(449):77–84, 2000.

S Lee and V Honavar. m-Transportability: Transportability of a Causal Effect from Multiple Environments. In *AAAI*. Citeseer, 2013.

A R Luedtke and M J van der Laan. Statistical inference for the mean outcome under a possibly non-unique optimal treatment strategy. *Annals of Statistics*, 44(2):713–742, 2016a.

A R Luedtke and M J van der Laan. Optimal Individualized Treatments in Resource-Limited Settings. *The international journal of biostatistics*, 12(1): 283–303, 2016b.

R F MacLehose, S Kaufman, J S Kaufman, and C Poole. Bounding causal effects under uncontrolled confounding using counterfactuals. *Epidemiology*, 16(4):548–555, 2005.

C F Manski. Nonparametric bounds on treatment effects. *The American Economic Review*, pages 319–323, 1990.

C F Manski. *Partial identification of probability distributions*. Springer, Berlin Heidelberg New York, 2003.

J Nealon, A-F Taurel, M R Capeding, N H Tran, S R Hadinegoro, T Chotpitayasunondh, C K Chong, T A Wartel, S Beucher, C Frago, and Others. Symptomatic Dengue Disease in Five Southeast Asian Countries: Epidemiological Evidence from a Dengue Vaccine Trial. *PLoS Negl Trop Dis*, 10(8): e0004918, 2016.

J Pearl and E Bareinboim. External validity: From do-calculus to transportability across populations. *Statistical Science*, 29(4):579–595, 2014.

E Polley and M J van der Laan. SuperLearner: super learner prediction, 2013. URL http://cran.r-project.org/package=SuperLearner.

R Core Team. R: a language and environment for statistical computing, 2014. URL http://www.r-project.org/.

J M Robins. A new approach to causal inference in mortality studies with a sustained exposure periodapplication to control of the healthy worker survivor effect. *Mathematical Modelling*, 7(9):1393–1512, 1986.

J M Robins. Optimal structural nested models for optimal sequential decisions. In D Y Lin and Heagerty P, editors, *Proceedings of the Second Seattle Symposium in Biostatistics*, volume 179, pages 189–326, 2004.

J M Robins and A Rotnitzky. Semiparametric Efficiency in Multivariate Regression Models with Missing Data. *J Am Stat Assoc*, 90:122–129, 1995.

S Rose and M J van der Laan. A Targeted Maximum Likelihood Estimator for Two-Stage Designs. *Int J Biostat*, 7(17), 2011.

D B Rubin and M J van der Laan. Empirical efficiency maximization: improved locally efficient covariate adjustment in randomized experiments and survival analysis. *Int J Biostat*, 4(1):Article 5, 2008.

K E Rudolph and M J van der Laan. Robust estimation of encouragement design intervention effects transported across sites. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2016.

A Sabchareon, D Wallace, C Sirivichayakul, K Limkittikul, P Chanthavanich, S Suvannadabba, V Jiwariyavej, W Dulyachai, K Pengsaa, T A Wartel, and Others. Protective efficacy of the recombinant, live-attenuated, CYD tetravalent dengue vaccine in Thai schoolchildren: a randomised, controlled phase 2b trial. *The Lancet*, 380(9853):1559–1567, 2012.

E Sarti, M L'Azou, M Mercado, P Kuri, J B Siqueira, E Solis, F Noriega, and R L Ochiai. A comparative study on active and passive epidemiological surveillance for dengue in five countries of Latin America. *International Journal of Infectious Diseases*, 44:44–49, 2016.

M J van der Laan and S Gruber. One-Step Targeted Minimum Loss-based Estimation Based on Universal Least Favorable One-Dimensional Submodels. *The international journal of biostatistics*, 12(1):351–378, 2016.

M J van der Laan and J M Robins. *Unified methods for censored longitudinal data and causality*. Springer, New York Berlin Heidelberg, 2003.

M J van der Laan and S Rose. *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer, New York, New York, 2011.

M J van der Laan, E Polley, and A Hubbard. Super Learner. *Stat Appl Genet Mol*, 6(1):Article 25, 2007. ISSN 1.

A W van der Vaart and J A Wellner. *Weak convergence and empirical processes*. Springer, Berlin Heidelberg New York, 1996.

S Vansteelandt, E Goetghebeur, M G Kenward, and G Molenberghs. Ignorance and uncertainty regions as inferential tools in a sensitivity analysis. *Statistica Sinica*, 16(3):953–979, 2006.

L Villar, Gustavo H Dayan, J L Arredondo-García, D M Rivera, R Cunha, C Deseda, H Reynales, M S Costa, J O Morales-Ramírez, G Carrasquilla, and Others. Efficacy of a tetravalent dengue vaccine in children in Latin America. *New England Journal of Medicine*, 372(2):113–123, 2015.

SAGE Working Group on Dengue Vaccines WHO. Background Paper on Dengue Vaccines. Technical report, World Health Organization, 2016. URL http://www.who.int/immunization/sage/meetings/2016/april.

Appendix

A Additional regularity conditions for Theorem 3

For general $\mathcal{P}' \in \mathcal{M}$, below we write VE'_- and $D'_{\mathrm{VE},s}$ for the analogues of $w \mapsto \mathrm{VE}^0_-(w)$ and $D^0_{\mathrm{VE},s}$ under $\mathcal{P}'$ rather than $\mathcal{P}^0$. Define

$$
\begin{aligned}
&\mathrm{Rem}_1(\mathcal{P}', \mathcal{P}^0)(w) \\
&\triangleq \sum_{s=1}^{S} \left| \mathbb{E}'_s[Y|A=0,w] - \mathbb{E}^0_s[Y|A=0,w] + \mathbb{E}^0_s\left[\frac{\mathbb{1}_{\{A=0\}}}{P'_s(A|w)} \left(Y - \mathbb{E}^0_s[Y|A,w]\right) \middle| w \right] \right| \\
&= \sum_{s=1}^{S} \left| \left(1 - \frac{P^0_s(A=0|w)}{P'_s(A=0|w)}\right) \left(\mathbb{E}'_s[Y|A=0,w] - \mathbb{E}^0_s[Y|A=0,w]\right) \right|, \\
&\mathrm{Rem}_2(\mathcal{P}', \mathcal{P}^0)(w) \\
&\triangleq c(w) - \mathrm{VE}^0_-(w) + \sum_{s=1}^{S} \mathbb{E}^0_s\left[D'_{\mathrm{VE},s}(O) \middle| w \right] \\
&= \left[\mathrm{VE}'_-(w) - \mathrm{VE}^0_-(w)\right] \left[1 - \frac{\mathbb{v}_0(w) d(0,w) + \sum_{s=1}^{S} \mathbb{v}_s(w)\, \mathbb{E}^0_s\left[Y|A=0,w\right]}{\mathbb{v}_0(w) d(0,w) + \sum_{s=1}^{S} \mathbb{v}_s(w)\, \mathbb{E}'_s\left[Y|A=0,w\right]} \right] \\
&\quad - \frac{\mathrm{VE}'_-(w) \sum_s \mathbb{v}_s(w) \left[1 - \frac{P^0_s(A=0|w)}{P'_s(A=0|w)}\right] \left(\mathbb{E}'_s[Y|A=0,w] - \mathbb{E}^0_s[Y|A=0,w]\right)}{\mathbb{v}_0(w) d(0,w) + \sum_{s=1}^{S} \mathbb{v}_s(w)\, \mathbb{E}^0_s\left[Y|A=0,w\right]}.
\end{aligned}
$$

We now state the additional regularity conditions needed for Theorem 3.

C2) $\mathcal{P}^n$ is consistent in the sense that all of the following terms are $o_P(n_{\min}^{-1/2})$:

$$
\begin{aligned}
&\left\| \mathrm{Rem}_1(\mathcal{P}^n, \mathcal{P}^0) \right\|_{2, P^0_\star}, \\
&\left\| \mathrm{Rem}_2(\mathcal{P}^n, \mathcal{P}^0) \right\|_{2, P^0_\star}.
\end{aligned}
$$

Furthermore, for each $s = 1, \ldots, S$, all of the following terms are also $o_P(n_{\min}^{-1/2})$:

$$
\begin{aligned}
&\sup_{\beta: \mathcal{W} \to [0,1]} \left\| w \mapsto \mathbb{E}^0_s[D^{\beta,n}_{\mathrm{UR},s}|w] - \mathbb{E}^0_s[D^{\beta,0}_{\mathrm{UR},s}|w] \right\|_{2,P^0_s} \left\| \frac{dP^n_\star}{dP^n_s} - \frac{dP^0_\star}{dP^0_s} \right\|_{2,P^0_s}, \\
&\left\| \frac{dP^n_\star}{dP^n_s} - \frac{dP^0_\star}{dP^0_s} \right\|_{2,P^0_s} \sup_{\beta: \mathcal{W} \to [0,1]} \left\| w \mapsto \mathbb{E}^0_s[D^{\beta,n}_{\mathrm{UR},s}(O)|w] - \mathbb{E}^0_s[D^{\beta,0}_{\mathrm{UR},s}(O)|w] \right\|_{2,P^0_s},
\end{aligned}
$$

$$\left\| \frac{dP^n_\star}{dP^n_s} - \frac{dP^0_\star}{dP^0_s} \right\|_{2,P^0_s} \left\| w \mapsto \mathbb{E}^0_s[D^n_{\mathrm{VE},s}(O)|w] - \mathbb{E}^0_s[D^0_{\mathrm{VE},s}(O)|w] \right\|_{2,P^0_s},$$
$$\left\| \upsilon^n - \upsilon^0 \right\|_{2,P^0_s} \left\| \mathrm{VE}^n - \mathrm{VE}^0_- \right\|_{2,P^0_s}.$$

Like Donsker conditions used in one-sample problems, C3 requires certain centered and scaled empirical means to be tight random elements. The next assumption makes this requirement formal.

C3) $(Q^n_s - P^0_s)\nabla\Omega^{\beta^0,n}_s = O_P(n^{-1/2}_{\min})$, $(Q^n_s - P^0_s)\nabla\Omega^{\widehat{\beta}^n,n}_s = O_P(n^{-1/2}_{\min})$ for each $s \in \mathcal{S}$.

In our multiple sample problem, one can analyze $(Q^n_s - P^0_s)\nabla\Omega^{\beta,n}_s$ for β fixed ($\beta = \beta^0$) or random ($\beta = \widehat{\beta}^n$) for each s by conditioning on all observations not belonging to trial s, and then applying a maximal inequality (which relies on the bracketing or uniform entropy integral of the class to which $\nabla\Omega^{\beta,n}_s$ belongs, see van der Vaart and Wellner, 1996) to bound the randomness in this term resulting from observations in trial s.

We also require a condition that is similar to asymptotic equicontinuity conditions implied by the use of Donsker classes in one-sample problems.

C4) $(Q^n_s - P^0_s)[\nabla\Omega^{\widehat{\beta}^n,n}_s - \nabla\Omega^{\beta^0,0}_s]$, $(Q^n_s - P^0_s)[\nabla\Gamma^{\widehat{\beta}^n,n}_s - \nabla\Gamma^{\beta^0,0}_s]$ for each $s \in \mathcal{S}$.

Using the same arguments outlined following C3, one can show that, if $\nabla\Omega^{\widehat{\beta}^n,n}_s$ and $\nabla\Gamma^{\widehat{\beta}^n,n}_s$ respectively converge to $\nabla\Omega^{\beta^0,0}_s$ and $\nabla\Gamma^{\beta^0,0}_s$ in P^0_s mean-square, we will allow be able to establish C4. While we do not formally give conditions under which this mean-square convergence occurs, the primary assumption needed in addition to the consistency conditions used in C2 is that $\widehat{\theta}^n$ is consistent for θ^0. We give sufficient conditions for this convergence in Lemma A.14.

The next assumption is made for simplicity, though it is very mild, especially under B4, which states that $w \mapsto \mathrm{VE}^0_-(w)$ is bounded away from $-\infty$.

C5) $\inf_w \mathrm{VE}^n(w) > -c > -\infty$ with probability approaching one, where the constant $c < \infty$ does not depend on sample size.

We make two additional assumptions when C1 holds. We first assume that the empirical version of $\omega^{\beta^0,0}$ is also close to μ, which seems reasonable given that $\omega^{\beta^0,0} = \mu$ under C1.

C6) $\widehat{\omega}^{\widehat{\beta}^n,n} - \mu = o_P(n^{-1/2}_{\min})$.

This condition can formally be supported in practice by looking at the size of $\widehat{\omega}^{\widehat{\beta}^n,n} - \mu$. The next regularity condition ensures that $\widehat{\beta}^n$ is a sufficiently good estimate of β^0 so that $\mathrm{UR}^{\widehat{\beta}^n,0}$ can be interpreted as being approximately equal to the the worst-case conditional unvaccinated risk.

C7) $\mathrm{Rem}_3^n \triangleq \gamma^{\widehat{\beta}^n,0} - \gamma^{\beta^0,0} - \theta^0[\omega^{\widehat{\beta}^n,0} - \omega^{\beta^0,0}] = o_P(n_{\min}^{-1/2})$.

We give more interpretable sufficient conditions for C7 in Appendix C. These conditions are a variant of the margin assumptions used in the classification literature.

B Proofs

B.1 Proof of lower bound

Proof of Lemma 1. Because $\mu = \mathbb{E}_\star^0[\mathbb{E}_\star^{0,F}[Y|A=0,W]]$ and B2 holds, $\omega^{\beta^0,0} = \mu$. By B1, $\omega^{\beta^0,0}\Psi(P_\star^{0,F}) \geq \mathbb{E}_\star^0[\mathbb{E}_\star^{0,F}[Y|A=0,W]\,\mathrm{VE}_-^0(W)]$. Below we respectively use $\mathcal{E}_1$, $\mathcal{E}_2$, and $\mathcal{E}_3$ to denote the events $\{\mathrm{VE}_-^0(W) < \theta^0\}$, $\{\mathrm{VE}_-^0(W) = \theta^0\}$, and $\{\mathrm{VE}_-^0(W) > \theta^0\}$, we see that

$$\begin{aligned}
&\omega^{\beta^0,0}\Psi(P_\star^{0,F}) - \gamma^{\beta^0,0} \geq \mathbb{E}_\star^0[\mathbb{E}_\star^{0,F}[Y|A=0,W]\,\mathrm{VE}_-^0(W)] - \gamma^{\beta^0,0} \\
&\quad = \mathbb{E}_\star^0\left[(\mathbb{1}_{\mathcal{E}_1} + \eta^0\,\mathbb{1}_{\mathcal{E}_2})\left\{\mathbb{E}_\star^{0,F}[Y|A=0,W] - \upsilon^0(W)\right\}\mathrm{VE}_-^0(W)\right] \\
&\qquad + \mathbb{E}_\star^0\left[(\mathbb{1}_{\mathcal{E}_3} + [1-\eta^0]\,\mathbb{1}_{\mathcal{E}_2})\left\{\mathbb{E}_\star^{0,F}[Y|A=0,W] - \lambda^0(W)\right\}\mathrm{VE}_-^0(W)\right].
\end{aligned}$$

By B2 and the fact that $\eta^0 \in [0,1]$, replacing each instance of $\mathrm{VE}_-^0(W)$ on the right-hand side by θ^0 yields a lower bound of the form $\theta^0\left[\mu - \omega^{\beta^0,0}\right]$. As $\omega^{\beta^0,0} = \mu$, this lower bound is equal to zero. Dividing both sides by $\omega^{\beta^0,0} > 0$ gives the result. □

Lemma A.1. *If $\lambda^0 \equiv 0$, B3, and B4, then ϕ_μ^0 is monotonically nondecreasing in μ.*

Proof of Lemma A.1. Fix $\theta_1 \leq \theta_2$ and $x_1, x_2 \in [0,1]$. If $\theta_1 = \theta_2$, then suppose also that $x_1 \leq x_2$. For $k = 1, 2$, define $\beta_k \triangleq w \mapsto \mathbb{1}_{\{\mathrm{VE}_-^0(w) < \theta_k\}} + x_k\,\mathbb{1}_{\{\mathrm{VE}_-^0(w) = \theta_k\}}$. Note that $\beta_2(w) - \beta_1(w) \geq 0$ and is strictly positive only if $\mathrm{VE}_-^0(W) \geq \theta_1$. Observe that

$$\frac{\gamma^{\beta_1,0}}{\omega^{\beta_1,0}} - \frac{\gamma^{\beta_2,0}}{\omega^{\beta_2,0}} = \frac{\left[\omega^{\beta_2,0} - \omega^{\beta_1,0}\right]\gamma^{\beta_1,0} - \omega^{\beta_1,0}\left[\gamma^{\beta_2,0} - \gamma^{\beta_1,0}\right]}{\omega^{\beta_1,0}\omega^{\beta_2,0}}$$

$$= \frac{\mathbb{E}^0_\star\left[\{\beta_2(W)-\beta_1(W)\}\upsilon^0(W)\{\gamma^{\beta_1,0}-\mathrm{VE}^0_-(W)\omega^{\beta_1,0}\}\right]}{\omega^{\beta_1,0}\omega^{\beta_2,0}}$$
$$\leq \frac{\mathbb{E}^0_\star\left[\{\beta_2(W)-\beta_1(W)\}\upsilon^0(W)\right]\left[\gamma^{\beta_1,0}-\theta_1\omega^{\beta_1,0}\right]}{\omega^{\beta_1,0}\omega^{\beta_2,0}}.$$

Noting that

$$\gamma^{\beta_1,0} = \mathbb{E}^0_\star[\upsilon^0(W)\{\mathbb{1}_{\{\mathrm{VE}^0_-(w)<\theta_1\}} + x_1\,\mathbb{1}_{\{\mathrm{VE}^0_-(w)=\theta_1\}}\}\,\mathrm{VE}^0_-(W)]$$
$$\leq \theta_1\,\mathbb{E}^0_\star[\upsilon^0(W)\{\mathbb{1}_{\{\mathrm{VE}^0_-(w)<\theta_1\}} + x_1\,\mathbb{1}_{\{\mathrm{VE}^0_-(w)=\theta_1\}}\}] = \theta_1\omega^{\beta_1,0},$$

we see that $\frac{\gamma^{\beta_1,0}}{\omega^{\beta_1,0}} \leq \frac{\gamma^{\beta_2,0}}{\omega^{\beta_2,0}}$.

We now write θ^0_μ and x^0_μ to make the dependence of θ^0, η^0 on μ explicit. As $\upsilon^0(w) > 0$ for all w, clearly $\mu_1 < \mu_2$ implies that $\theta^0_{\mu_1} \leq \theta^0_{\mu_2}$ and, if $\theta^0_{\mu_1} = \theta^0_{\mu_2}$, then also that $x^0_{\mu_1} \leq x^0_{\mu_2}$. This completes the proof. □

B.2 Pathwise derivative of Φ (Theorem 2)

Throughout this section, we refer to parameters evaluated at the collection of distributions $\mathcal{P}^\epsilon$ rather than parameters evaluated at $\mathcal{P}^0$ by replacing the superscript zero by superscript ϵ, e.g. we write $\upsilon^\epsilon \triangleq \Upsilon(\mathcal{P}^\epsilon)$ rather than $\upsilon^0 \triangleq \Upsilon(\mathcal{P}^0)$.

We now prove Theorem 2. The proof references results that we prove later in this section.

Proof of Theorem 2. We consider the cases that C1 holds and that C1' holds separately. The proof in the case where C1" holds is nearly identical that under C1', and so the proof is omitted.

Case 1: C1 holds. By C1 and the fact that υ^0 is bounded away from zero, $\omega^{\beta^0,0} = \mu$ and θ^0 is finite. By Lemma A.6, θ^ϵ is finite for all ϵ small enough so that $\omega^{\beta^\epsilon,\epsilon} = \mu$ for these ϵ. Hence, $\phi^\epsilon - \phi^0 = \frac{\gamma^{\beta^\epsilon,\epsilon} - \gamma^{\beta^0,0}}{\mu}$ for all ϵ sufficiently small. Dividing both sides by ϵ, taking the limit as $\epsilon \to 0$, and applying Theorem A.3 (Section B.2.2) shows that the gradients of Φ at $\mathcal{P}^0$ are given by the desired expressions.

Case 2: C1' holds. The key to proving Theorem 2 under C1' is proving that $P^\epsilon_\star\{\mathrm{VE}^\epsilon_-(W) < \theta^\epsilon\} = P^0_\star\{\mathrm{VE}^0_-(W) < \theta^0\} = 1$ for all ϵ small enough. We formally prove this in Lemma A.2 (Section B.2.1). It follows that

$$\Phi(\mathcal{P}^\epsilon) - \Phi(\mathcal{P}^0) = \frac{\mathbb{E}^\epsilon_\star\left[\upsilon^\epsilon(W)\,\mathrm{VE}^\epsilon(W)\right]}{\mathbb{E}^\epsilon_\star\left[\upsilon^\epsilon(W)\right]} - \frac{\mathbb{E}^0_\star\left[\upsilon^0(W)\,\mathrm{VE}^0_-(W)\right]}{\mathbb{E}^0_\star\left[\upsilon^0(W)\right]}.$$

Dividing both sides by ϵ and taking the limit as $\epsilon \to 0$ yields the same expression that would be used to evaluate the gradients of the parameter $\mathcal{P}' \to \frac{\int \Upsilon(\mathcal{P}')(w)\, \mathrm{VE}'_{-}(w) dP'_\star(w)}{\int \Upsilon(\mathcal{P}')(w) dP'_\star(w)}$. It is straightforward to derive expressions for the gradients of this parameter via the delta method. □

B.2.1 Lemma used to Prove Theorem 2 under C1'.

Lemma A.2. *If B3, B4, and C1', then* $P^\epsilon_\star\{\mathrm{VE}^\epsilon_{-}(W) < \theta^\epsilon\} = 1$ *for all* ϵ *small enough.*

Proof of Lemma A.2. Standard pathwise differentiability calculations show that

$$\begin{aligned}\frac{\mathbb{E}^\epsilon_\star[\upsilon^\epsilon(W)] - \omega^{\beta^0,0}}{\epsilon} &= \frac{\mathbb{E}^\epsilon_\star[\upsilon^\epsilon(W)] - \mathbb{E}^0_\star[\upsilon^0(W)]}{\epsilon} \\ &= \sum_{s\in\mathcal{S}} P^\epsilon_s \nabla\Omega^{\beta^0,0}_s + o(1),\end{aligned}$$

and the fact that each $\nabla\Omega^{\beta^0,0}_s(O_s)$ has mean zero for $O_s \sim P^0_s$ implies that the right-hand side is $O(1)$. By C1', $\omega^{\beta^0,0} < \mu$. Hence, $\mathbb{E}^\epsilon_\star[\upsilon^\epsilon(W)] < \mu + O(\epsilon)$. It follows that, for all ϵ sufficiently small, $\mathbb{E}^\epsilon_\star[\mathrm{UR}^{w\mapsto \mathbb{1}_{\{\mathrm{VE}^\epsilon_{-}(w)<\theta\},\epsilon}}(W)] < \mu$ for all θ. Hence, $\theta^\epsilon = +\infty$ for all ϵ sufficiently small, completing the proof. □

B.2.2 Theorem used to prove Theorem 2 under C1.

The following theorem establishes the pathwise differentiability of the parameter $\mathcal{P}' \mapsto \Gamma^{\beta'}(\mathcal{P}')$, where the β' in the subscript is the analogue of β^0 but defined at parameter input $\mathcal{P}'$ rather than at $\mathcal{P}^0$.

Theorem A.3. *If B3, B4, and C1, then* $\mathcal{P}' \mapsto \Gamma^{\beta'}(\mathcal{P}')$ *is pathwise differentiable at* $\mathcal{P}^0$ *with* P^0_s *gradients* $o_s \mapsto \nabla\Gamma^{\beta^0,0}_s(o_s) - \theta^0\nabla\Omega^{\beta^0,0}_s(o_s)$, $s \in \mathcal{S}$.

Proof of Theorem A.3. Observe that

$$\frac{\gamma^{\beta^\epsilon,\epsilon} - \gamma^{\beta^0,0}}{\epsilon} = \frac{\gamma^{\beta^0,\epsilon} - \gamma^{\beta^0,0}}{\epsilon} + \frac{\gamma^{\beta^\epsilon,\epsilon} - \gamma^{\beta^0,\epsilon}}{\epsilon}.$$

We wish to obtain an expression for the limit on the left as $\epsilon \to 0$. If the limits exist for both terms on the right, then the limit on the left is given by their sum. A standard delta method argument shows that the limit as $\epsilon \to 0$ on the right is given by $\sum_{s\in\mathcal{S}} P^0_s \nabla\Gamma^{\beta^0,0}_s h_s$. Lemma A.4 shows that the limit exists for the latter term and provides its closed form expression. □

Lemma A.4. *Under the conditions of Theorem A.3,*

$$\lim_{\epsilon \to 0} \frac{\gamma^{\beta^\epsilon,\epsilon} - \gamma^{\beta^0,\epsilon}}{\epsilon} = -\theta^0 \sum_{s \in \mathcal{S}} \int \nabla \Omega_s^{\beta^0,0}(o_s) h_s(o_s) dP_s^0(o_s).$$

Proof of Lemma A.4. Telescoping yields that

$$\begin{aligned}
\gamma^{\beta^\epsilon,\epsilon} - \gamma^{\beta^0,\epsilon} &= \mathbb{E}_\star^\epsilon \left[\left\{ \mathrm{UR}^{\beta^\epsilon,\epsilon}(W) - \mathrm{UR}^{\beta^0,\epsilon}(W) \right\} \mathrm{VE}^\epsilon(W) \right] \\
&= \mathbb{E}_\star^\epsilon \left[\left\{ \mathrm{UR}^{\beta^\epsilon,\epsilon}(W) - \mathrm{UR}^{\beta^0,\epsilon}(W) \right\} \left\{ \mathrm{VE}^\epsilon(W) - \theta^0 \right\} \right] \\
&\quad + \theta^0 \int \left\{ \mathrm{UR}^{\beta^\epsilon,\epsilon}(w) dP_\star^\epsilon(w) - \mathrm{UR}^{\beta^0,0}(w) dP_\star^0(w) \right\} \\
&\quad - \theta^0 \int \left\{ \mathrm{UR}^{\beta^0,\epsilon}(w) dP_\star^\epsilon(w) - \mathrm{UR}^{\beta^0,0}(w) dP_\star^0(w) \right\}.
\end{aligned}$$

We denote the first and second terms in the final equality by T_1^ϵ and T_2^ϵ. If we can show that T_1^ϵ and T_2^ϵ are $o(\epsilon)$, then the result is immediate by dividing both sides by ϵ and applying the chain rule to the third term.

The remainder of this proof uses a positive constant c that may vary line by line. We first establish that $T_1^\epsilon = o(\epsilon)$. Straightforward calculations show that

$$\begin{aligned}
|T_1^\epsilon| \leq (1 + c|\epsilon|) \, \mathbb{E}_\star^0 \Bigg[& \left\{ \left| \mathrm{UR}^{\beta^\epsilon,0}(W) - \mathrm{UR}^{\beta^0,0}(W) \right| + c|\epsilon| \right\} \\
& \times \left\{ \left| \mathrm{VE}_-^0(W) - \theta^0 \right| + c|\epsilon| \right\} \Bigg].
\end{aligned}$$

For simplicity we suppose that $\eta^\epsilon = 0$ for the remainder of the proof, but the proof with $\eta^\epsilon \neq 0$ is nearly identical. For each w,

$$\left| \mathrm{UR}^{\beta^\epsilon,0}(w) - \mathrm{UR}^{\beta^0,0}(w) \right| = \left| \mathbb{1}_{\{\mathrm{VE}_-^\epsilon(w) < \theta^\epsilon\}} - \mathbb{1}_{\{\mathrm{VE}_-^0(w) < \theta^0\}} \right| [\upsilon^0(w) - \lambda^0(w)].$$

Using that $\mathbb{1}_{\{\mathrm{VE}_-^\epsilon(w) < \theta^\epsilon\}} \neq \mathbb{1}_{\{\mathrm{VE}_-^0(w) < \theta^0\}}$ implies that $0 \leq |\mathrm{VE}_-^0(w) - \theta^0| < |\mathrm{VE}_-^\epsilon(w) - \theta^\epsilon - \mathrm{VE}_-^0(w) + \theta^0|$ then shows that

$$\begin{aligned}
|T_1^\epsilon| \leq (1 + c|\epsilon|) \, \mathbb{E}_\star^0 \Bigg| & \mathbb{1}_{\{0 \leq |\mathrm{VE}_-^0(w) - \theta^0| < |\mathrm{VE}_-^\epsilon(w) - \theta^\epsilon - \mathrm{VE}_-^0(w) + \theta^0|\}} \\
& \times \left\{ [\upsilon^0(W) - \lambda^0(W)] |\mathrm{VE}_-^0(W) - \theta^0| + c|\epsilon| \right\} \Bigg|.
\end{aligned}$$

The indicator above can be replaced by $\mathbb{1}_{\{0<|\operatorname{VE}^0_-(w)-\theta^0|<|\operatorname{VE}^\epsilon_-(w)-\theta^\epsilon-\operatorname{VE}^0_-(w)+\theta^0|\}}$ because C1 implies that $P^0_\star(\operatorname{VE}^0_-(W)=\theta^0)=0$. Combining this with the fact that $v^0(W)-\lambda^0(W)\leq 1$ and $|\operatorname{VE}^\epsilon_-(w)-\operatorname{VE}^0_-(w)|\leq c|\epsilon|$ shows that

$$|T^\epsilon_1|\leq (1+c|\epsilon|)\,\mathbb{E}^0_\star\left[\mathbb{1}_{\{0<|\operatorname{VE}^0_-(w)-\theta^0|<|\operatorname{VE}^\epsilon_-(w)-\theta^\epsilon-\operatorname{VE}^0_-(w)+\theta^0|\}}\left\{|\theta^\epsilon-\theta^0|+c|\epsilon|\right\}\right].$$

By Lemma A.6, $\theta^\epsilon-\theta^0=O(\epsilon)$ so that, for all ϵ large enough,

$$|T^\epsilon_1|\leq c|\epsilon|P^0_\star\left\{0<|\operatorname{VE}^0_-(w)-\theta^0|<c|\epsilon|\right\}.$$

The probability statement is $o(1)$. Thus, $T^\epsilon_1=o(\epsilon)$.

Lemma A.6 establishes that θ^ϵ is finite for all ϵ small enough, and

The assumption that C1 establishes that θ^0 is finite, and thus Lemma A.6 establishes that θ^ϵ is finite for all ϵ small enough. It follows that $\omega^{\beta^\epsilon,\epsilon}=\mu$ for all ϵ small enough, and thus $T^\epsilon_2=\omega^{\beta^\epsilon,\epsilon}-\omega^{\beta^0,0}=0$ for these ϵ. □

Lemma A.5. *For each t, let $\tilde\beta_t\triangleq w\mapsto\mathbb{1}_{\{\operatorname{VE}^0_-(w)<t\}}$. If C1 holds, then, for all θ in a neighborhood of θ^0,*

$$0<\liminf_{t\to 0}\frac{\omega^{\tilde\beta_{\theta+t},0}-\omega^{\tilde\beta_\theta,0}}{t}\leq\limsup_{t\to 0}\frac{\omega^{\tilde\beta_{\theta+t},0}-\omega^{\tilde\beta_\theta,0}}{t}<\infty,$$

Proof of Lemma A.5. As $v^0(w)-\lambda^0(w)\geq\delta$ for all w, the facts that $t\mapsto P^0_\star\{\operatorname{VE}^0_-(W)<\theta^0+t\}$ is monotonically non-decreasing and that C1 holds yield that

$$\begin{aligned}\liminf_{t\to 0}\frac{\omega^{\tilde\beta_{\theta+t},0}-\omega^{\tilde\beta_\theta,0}}{t}&=\liminf_{t\to 0}\frac{\mathbb{E}^0_\star[\{v^0(W)-\lambda^0(W)\}I\{\theta^0\leq\operatorname{VE}^0_-(W)<\theta^0+t\}]}{t}\\&\geq\delta\liminf_{t\to 0}\frac{P^0_\star\{\operatorname{VE}^0_-(W)<\theta^0+t\}-P^0_\star\{\operatorname{VE}^0_-(W)<\theta^0\}}{t}>0.\end{aligned}$$

The limit supremum result is proven similarly. □

Lemma A.6. *Under the conditions of Theorem A.3, $\theta^\epsilon=\theta^0+O(\epsilon)$.*

Proof of Lemma A.6. In what follows we abuse notation and write $\omega^{\theta,\epsilon}$ for $\omega^{w\mapsto\mathbb{1}_{\{\operatorname{VE}^\epsilon_-(w)<\theta\},\epsilon}}$ for all θ,ϵ. Note that, by C1, $\eta^0=0$ so that $\omega^{\beta^0,0}=\omega^{\theta^0,0}$. We similarly abuse notation by letting $\operatorname{UR}^{\theta,\epsilon}\triangleq\operatorname{UR}^{w\mapsto\mathbb{1}_{\{\operatorname{VE}^\epsilon_-(w)<\theta\},\epsilon}}$ and $D^{\theta,\epsilon}_{\operatorname{UR},s}\triangleq D^{w\mapsto\mathbb{1}_{\{\operatorname{VE}^\epsilon_-(w)<\theta\},\epsilon}}_{\operatorname{UR},s}$.

We first establish that $\theta^\epsilon \to \theta^0$. For $t > 0$, the definition of θ^ϵ yields that $|\theta^\epsilon - \theta^0| > t$ implies that either $\omega^{\theta^0 - t,\epsilon} > \mu$ or $\omega^{\theta^0+t,\epsilon} \leq \mu$. We will establish that $\theta^\epsilon \to \theta^0$ by showing that, for any $t > 0$, both (i) $\omega^{\theta^0-t,\epsilon} \leq \mu$ and (ii) $\omega^{\theta^0+t,\epsilon} > \mu$, and therefore $|\theta^\epsilon - \theta^0| \leq t$.

For a constant $c > 0$ that may vary line by line,

$$\begin{aligned}
\omega^{\theta^0-t,\epsilon} &\leq (1 + c|\epsilon|)\, \mathbb{E}^0_\star \left[\mathrm{UR}^{\theta^0-t,\epsilon}(W)\right] \\
&= (1 + c|\epsilon|)\, \mathbb{E}^0_\star \left[\lambda^\epsilon(W) + \{\upsilon^\epsilon(W) - \lambda^\epsilon(W)\}\, \mathbb{1}\{\mathrm{VE}^\epsilon_-(W) < \theta^0 - t\}\right] \\
&\leq (1 + c|\epsilon|)\, \mathbb{E}^0_\star \left[\lambda^\epsilon(W) + \{\upsilon^\epsilon(W) - \lambda^\epsilon(W)\}\, \mathbb{1}\{\mathrm{VE}^0_-(W) < \theta^0 - t + c|\epsilon|\}\right] \\
&= (1 + c|\epsilon|)\omega^{\theta^0-t+c|\epsilon|,0} + c|\epsilon|.
\end{aligned}$$

Now, we know that $\omega^{\theta^0-t+c|\epsilon|,0} \leq \mu$ by the definition of the supremum, and by C1 the inequality is strict. It follows that $\omega^{\theta^0-t,\epsilon} < \mu + c|\epsilon|(1 + \mu)$, and so $\omega^{\theta^0-t,\epsilon} < \mu$ for all $|\epsilon|$ sufficiently small. Thus (i) holds. We now establish (ii). Similarly to the above, we have that $\omega^{\theta^0+t,\epsilon} \geq (1 - c|\epsilon|)\omega^{\theta^0+t-c|\epsilon|,0} - c|\epsilon|$ for a constant $c > 0$. By the definition of the supremum, $\omega^{\theta^0+t-c|\epsilon|,0} > \mu$ for all $|\epsilon|$ sufficiently small. Thus, (ii) holds for all $|\epsilon|$ sufficiently small. From the observation at the beginning of the proof, the fact that (i) and (ii) hold for all $|\epsilon|$ sufficiently small implies that $\limsup_{\epsilon\to 0} |\theta^\epsilon - \theta^0| \leq t$, and as t was arbitrary this implies that $\theta^\epsilon - \theta^0 = o(1)$.

We now establish that $\theta^\epsilon - \theta^0 = O(\epsilon)$. Using similar techniques to those used above, one can show that there exists a $c > 0$ such that, for all θ,

$$(1 - c|\epsilon|)\omega^{\theta-c|\epsilon|,0} - c|\epsilon| \leq \omega^{\theta,\epsilon} \leq (1 + c|\epsilon|)\omega^{\theta+c|\epsilon|,0} + c|\epsilon|. \tag{A.1}$$

Noting that $\mu \geq \omega^{\theta^\epsilon,\epsilon}$, we have that

$$\begin{aligned}
0 &\geq \omega^{\theta^\epsilon,\epsilon} - \omega^{\theta^0,0} \geq (1 - c|\epsilon|)\omega^{\theta^\epsilon-c|\epsilon|,0} - c|\epsilon| - \omega^{\theta^0,0} \\
&= (1 - c|\epsilon|) \left[\omega^{\theta^\epsilon-c|\epsilon|,0} - \omega^{\theta^0,0}\right] - c|\epsilon|(1 + \omega^{\theta^0,0}).
\end{aligned}$$

As C1 implies $\liminf_{t\to 0} \frac{\omega^{\theta^0+t,0} - \omega^{\theta^0,0}}{t} > 0$ by Lemma A.5, $B \triangleq \liminf_{\epsilon\to 0} B^\epsilon \triangleq \liminf_{\epsilon\to 0} \frac{\omega^{\theta^\epsilon-c|\epsilon|,0} - \omega^{\theta^0,0}}{\theta^\epsilon - \theta^0 - c|\epsilon|}$ is greater than zero. The above implies that, for all ϵ small enough,

$$\frac{c|\epsilon|(1 + \omega^{\theta^0,0})}{1 - c|\epsilon|} \geq \omega^{\theta^\epsilon-c|\epsilon|,0} - \omega^{\theta^0,0} = (\theta^\epsilon - \theta^0 - c|\epsilon|)B^\epsilon.$$

Dividing both sides by B^ϵ and taking the limit supremum of the left-hand side as $\epsilon \to 0$ shows that $\theta^\epsilon - \theta^0 \leq O(\epsilon)$.

The proof of the fact that $\theta^\epsilon - \theta^0 \geq O(\epsilon)$ is analogous, making use of the upper bound in (A.1) and the fact that $\mu < \omega^{\theta^\epsilon+c|\epsilon|,\epsilon}$ if θ^ϵ is finite (guaranteed for all ϵ small enough). □

B.3 Asymptotic linearity (Theorem 3)

In what follows, we write C1/C6/C7 to mean that C1, C6, and C7 hold. The proof under C1" is essentially equivalent to that under C1', so for simplicity we only give the proof under C1' and C1/C6/C7. Outside of formal theorem statements, when we write that C1/C6/C7 holds or that C1' holds, we mean that the stated condition(s) and the other conditions of Theorem 3, namely B3, B4, C2, C3, C4, and C5, hold.

We prove Theorem 3 by analyzing the terms on the right-hand side of the following decomposition, whose derivation is straightforward:

$$\begin{aligned}
\widehat{\phi}^n - \phi^0 =& \frac{\widehat{\gamma}^{\widehat{\beta}^n,n}}{\widehat{\omega}^{\widehat{\beta}^n,n}} - \frac{\gamma^{\beta^0,0}}{\omega^{\beta^0,0}} \\
=& \frac{\widehat{\gamma}^{\widehat{\beta}^n,n} - \gamma^{\widehat{\beta}^n,0}}{\omega^{\beta^0,0}} - \frac{\gamma^{\beta^0,0}}{\left[\omega^{\beta^0,0}\right]^2}\left[\widehat{\omega}^{\widehat{\beta}^n,n} - \omega^{\beta^0,0}\right] + \frac{\gamma^{\widehat{\beta}^n,0} - \gamma^{\beta^0,0}}{\omega^{\beta^0,0}} \\
&+ \left[\widehat{\gamma}^{\widehat{\beta}^n,n} - \gamma^{\beta^0,0}\right]\left[\frac{1}{\widehat{\omega}^{\widehat{\beta}^n,n}} - \frac{1}{\omega^{\beta^0,0}}\right] + \frac{\gamma^{\beta^0,0}}{\widehat{\omega}^{\widehat{\beta}^n,n}\left[\omega^{\beta^0,0}\right]^2}\left[\widehat{\omega}^{\widehat{\beta}^n,n} - \omega^{\beta^0,0}\right]^2 \\
\triangleq& \text{ Term 1} - \text{Term 2} + \text{Term 3} + \text{Term 4} + \text{Term 5}. \qquad \text{(A.2)}
\end{aligned}$$

We will show that Term 1 is non-negligible, i.e. contributes to the dominant $O_P(n_{\min}^{-1/2})$ term in the asymptotically linear expansion, under both C1/C6/C7 and C1', Term 2 is negligible, i.e. $o_P(n_{\min}^{-1/2})$, under C1/C6/C7 and non-negligible under C1', and Term 3 is non-negligible under C1/C6/C7 and negligible under C1'. Terms 4 and 5 are remainder terms that one can ignore if they invoke the delta method. Indeed, if we show that $\widehat{\omega}^{\widehat{\beta}^n,n} - \omega^{\beta^0,0}$ and $\widehat{\gamma}^{\widehat{\beta}^n,n} - \gamma^{\beta^0,0}$ are $O_P(n_{\min}^{-1/2})$, then Terms 4 and 5 are negligible. As this is what we will show under both C1/C6/C7 and C1', Terms 4 and 5 can be shown to be negligible by our analysis of the first three terms.

The following theorem establishes the behavior of Term 1 by studying $\widehat{\gamma}^{\widehat{\beta}^n,n} - \gamma^{\widehat{\beta}^n,0}$.

Theorem A.7 (Term 1)**.** *If B4, C2, C4, and C5, then*

$$\widehat{\gamma}^{\widehat{\beta}^n,n} - \gamma^{\widehat{\beta}^n,0} = \sum_{s\in\mathcal{S}}(Q_s^n - P_s^0)\nabla\Gamma_s^{\beta^0,0} + o_P(n_{\min}^{-1/2}).$$

The proof of Theorem A.7 is given in Appendix B.3.1.

The fact that Term 2 is negligible under C1/C6/C7 is an immediate consequence of C6. We now present a theorem that establishes the behavior of Term 3 under C1/C6/C7.

Theorem A.8 (Term 3 under C1/C6/C7)**.** *If B4, C1, C2, C4, C6, and C7, then*

$$\gamma^{\widehat{\beta}^n,0} - \gamma^{\beta^0,0} = -\theta^0 \sum_{s\in\mathcal{S}} (Q_s^n - P_s^0)\nabla\Omega_s^{\beta^0,0} + o_P(n_{\min}^{-1/2}).$$

The proof of Theorem A.8 is given in Appendix B.3.2.

We now present a lemma and a theorem that respectively establish the behavior of Terms 3 and 2 under C1'. Both results are proven in Appendix B.3.3.

Lemma A.9 (Term 3 under C1')**.** *If B4, C1', C2, and C3, then* $\widehat{\theta}^n = +\infty$*, and consequently* $Q_\star^n\{\mathrm{VE}_-^n(W) < \widehat{\theta}^n\} = 1$*, with probability approaching one.*

The proof of Lemma A.9 is given in Appendix B.3.3. By C1', $\theta^0 = +\infty$, so that the above shows that $\widehat{\theta}^n = \theta^0$, and thus that Term 2 is zero, with probability approaching one.

Theorem A.10 (Term 2 under C1')**.** *If B4, C1', C2, and C4, then*

$$\widehat{\omega}^{\widehat{\beta}^n,n} - \omega^{\beta^0,0} = \sum_{s\in\mathcal{S}} (Q_s^n - P_s^0)\nabla\Omega_s^{\beta^0,0} + o_P(n_{\min}^{-1/2}).$$

The proof of Theorem A.10 is given in Appendix B.3.3.

B.3.1 Term 1 under C1/C6/C7 and C1'. We give a lemma before proving Theorem A.7.

Lemma A.11. *Under the conditions of Theorem A.7,*

$$\gamma^{\widehat{\beta}^n,n} - \gamma^{\widehat{\beta}^n,0} + \sum_{s\in\mathcal{S}} P_s^0 \nabla\Gamma_s^{\widehat{\beta}^n,n} = o_P(n_{\min}^{-1/2}).$$

Proof of Lemma A.11. Note that

$$\begin{aligned}
&\gamma^{\widehat{\beta}^n,n} - \gamma^{\widehat{\beta}^n,0} + \sum_{s\in\mathcal{S}} P_s^0 \nabla\Gamma_s^{\widehat{\beta}^n,n} \\
&= \mathbb{E}_\star^0\left[\mathrm{UR}^{\widehat{\beta}^n,n}(W)\,\mathrm{VE}^n(W)\right] - \gamma^{\widehat{\beta}^n,0} + \sum_{s=1}^{S} P_s^0 \nabla\Gamma_s^{\widehat{\beta}^n,n}. \qquad \text{(A.3)}
\end{aligned}$$

Furthermore,

$$\sum_{s=1}^{S} P_s^0 \nabla\Gamma_s^{\widehat{\beta}^n,n}$$

$$
\begin{aligned}
&= \sum_{s=1}^{S} \mathbb{E}_s^0 \left[\frac{dP_\star^n}{dP_s^n}(W) \left\{ D_{\mathrm{UR},s}^{\widehat{\beta}^n,n}(O)\, \mathrm{VE}^n(W) + \mathrm{UR}^{\widehat{\beta}^n,n}(W) D_{\mathrm{VE},s}^n(O) \right\} \right] \\
&= \mathbb{E}_\star^0 \left[\sum_{s=1}^{S} \left\{ D_{\mathrm{UR},s}^{\widehat{\beta}^n,n}(O)\, \mathrm{VE}^n(W) + \mathrm{UR}^{\widehat{\beta}^n,n}(W) D_{\mathrm{VE},s}^n(O) \right\} \right] \qquad \text{(A.4)} \\
&\quad + \sum_{s=1}^{S} \mathbb{E}_s^0 \left[\left\{ \frac{dP_\star^n}{dP_s^n}(W) - \frac{dP_\star^0}{dP_s^0}(W) \right\} \left\{ D_{\mathrm{UR},s}^{\widehat{\beta}^n,n}(O)\, \mathrm{VE}^n(W) + \mathrm{UR}^{\widehat{\beta}^n,n}(W) D_{\mathrm{VE},s}^n(O) \right\} \right].
\end{aligned}
$$

The law of total expectation and the fact that $D_{\mathrm{UR},s}^{\widehat{\beta}^n,0}$ and $D_{\mathrm{VE},s}^0$ are mean zero when applied to a random variable O_s drawn from the conditional distribution $O_s|W$ under P_s^0 yield that the latter line is bounded above by

$$
\begin{aligned}
&\sum_{s=1}^{S} \mathbb{E}_s^0 \Bigg[\left\{ \frac{dP_\star^n}{dP_s^n}(W) - \frac{dP_\star^0}{dP_s^0}(W) \right\} \\
&\qquad \times \Big\{ \left(\mathbb{E}_s^0[D_{\mathrm{UR},s}^{\widehat{\beta}^n,n}(O)|W] - \mathbb{E}_s^0[D_{\mathrm{UR},s}^{\widehat{\beta}^n,0}(O)|W] \right) \mathrm{VE}^n(W) \\
&\qquad\quad + \mathrm{UR}^{\widehat{\beta}^n,n}(W) \left(\mathbb{E}_s^0[D_{\mathrm{VE},s}^n(O)|W] - \mathbb{E}_s^0[D_{\mathrm{VE},s}^0(O)|W] \right) \Big\} \Bigg].
\end{aligned}
$$

The above is $o_P(n_{\min}^{-1/2})$ by Cauchy-Schwarz, C2, C5, and the fact that $\mathrm{UR}^{\widehat{\beta}^n,n}$, VE^n have bounded range. Furthermore, (A.4) simplifies to

$$
\begin{aligned}
&\mathbb{E}_\star^0 \left[\sum_{s=1}^{S} \left\{ D_{\mathrm{UR},s}^{\widehat{\beta}^n,n}(O)\, \mathrm{VE}^n(W) + \mathrm{UR}^{\widehat{\beta}^n,n}(W) D_{\mathrm{VE},s}^n(O) \right\} \right] \\
&= \mathbb{E}_\star^0 \left[\left\{ [\mathrm{UR}^{\widehat{\beta}^n,0}(W) - \mathrm{UR}^{\widehat{\beta}^n,n}(W)]\, \mathrm{VE}^n(W) + \mathrm{UR}^{\widehat{\beta}^n,n}(W) [\mathrm{VE}_-^0(W) - \mathrm{VE}^n(W)] \right\} \right] \\
&\quad + \mathbb{E}_\star^0 \left[\left\{ \mathrm{UR}^{\widehat{\beta}^n,n}(W) - \mathrm{UR}^{\widehat{\beta}^n,0}(W) + \sum_{s=1}^{S} D_{\mathrm{UR},s}^{\widehat{\beta}^n,n}(O) \right\} \mathrm{VE}^n(W) \right] \\
&\quad + \mathbb{E}_\star^0 \left[\mathrm{UR}^{\widehat{\beta}^n,n}(W)\, \mathrm{Rem}_2(\mathcal{P}^n, \mathcal{P}^0) \right].
\end{aligned}
$$

The final line is $o_P(n_{\min}^{-1/2})$ by C2. The magnitude of the second line is upper bounded by $\mathbb{E}_\star^0 \left[\mathrm{Rem}_1(\mathcal{P}^n, \mathcal{P}^0)(W) |\, \mathrm{VE}^n(W)| \right]$, which is also $o_P(n_{\min}^{-1/2})$ by C2. We have thus established that (A.3) rewrites as

$$
\gamma^{\widehat{\beta}^n,n} - \gamma^{\widehat{\beta}^n,0} + \sum_{s \in \mathcal{S}} P_s^0 \nabla \Gamma_s^{\widehat{\beta}^n,n}
$$

$$= \mathbb{E}^0_\star \left[\{ \mathrm{UR}^{\widehat{\beta}^n,n}(W) - \mathrm{UR}^{\widehat{\beta}^n,0}(W) \} \{ \mathrm{VE}^0_-(W) - \mathrm{VE}^n(W) \} \right] + o_P(n_{\min}^{-1/2}).$$

By Cauchy-Schwarz and C2, the leading term on the right is $o_P(n_{\min}^{-1/2})$, so that indeed the entire right-hand side is $o_P(n_{\min}^{-1/2})$. □

We now prove Theorem A.7.

Proof of Theorem A.7. Rearranging the result of Lemma A.11 yields that

$$\gamma^{\widehat{\beta}^n,n} - \gamma^{\widehat{\beta}^n,0} = -\sum_{s \in \mathcal{S}} P_s^0 \nabla \Gamma_s^{\widehat{\beta}^n,n} + o_P(n_{\min}^{-1/2}).$$

Adding $\sum_{s \in \mathcal{S}} Q_s^n \nabla \Gamma_s^{\widehat{\beta}^n,n}$ to both sides shows that the one-step estimator $\widehat{\gamma}^{\widehat{\beta}^n,n}$ satisfies the identity

$$\widehat{\gamma}^{\widehat{\beta}^n,n} - \gamma^{\widehat{\beta}^n,0} = \sum_{s \in \mathcal{S}} (Q_s^n - P_s^0) \nabla \Gamma_s^{\widehat{\beta}^n,n} + o_P(n_{\min}^{-1/2}).$$

Applying C4 allows one to replace each instance of $\nabla \Gamma_s^{\widehat{\beta}^n,n}$ above by $\nabla \Gamma_s^{\beta^0,0}$, $s \in \mathcal{S}$. □

B.3.2 Term 3 under C1/C6/C7. We first give two lemmas, and then we establish control of Term 3 by proving Theorem A.8.

Lemma A.12. *If B4 and C2, then*

$$\left| \widehat{\omega}^{\widehat{\beta}^n,n} - \omega^{\widehat{\beta}^n,0} - \sum_{s \in \mathcal{S}} (Q_s^n - P_s^0) \nabla \Omega_s^{\widehat{\beta}^n,n} \right| = o_P(n_{\min}^{-1/2}).$$

Proof of Lemma A.12. Note that

$$\begin{aligned}
&\omega^{\widehat{\beta}^n,n} + \sum_{s \in \mathcal{S}} P_s^0 \nabla \Omega_s^{\widehat{\beta}^n,n} \\
&= \mathbb{E}^0_\star \left[\mathrm{UR}^{\widehat{\beta}^n,n}(W) \right] + \sum_{s=1}^S \mathbb{E}^0_s \left[\frac{dP_\star^n}{dP_s^n}(W) D_{\mathrm{UR},s}^{\widehat{\beta}^n,n}(O) \right] \\
&= \omega^{\widehat{\beta}^n,0} + \mathbb{E}^0_\star \left[\mathrm{UR}^{\widehat{\beta}^n,n}(W) - \mathrm{UR}^{\widehat{\beta}^n,0}(W) + \sum_{s=1}^S \mathbb{E}^0_s \left[D_{\mathrm{UR},s}^{\widehat{\beta}^n,n}(O) \middle| W \right] \right] \\
&\quad + \sum_{s=1}^S \mathbb{E}^0_s \left[\left\{ \frac{dP_\star^n}{dP_s^n}(W) - \frac{dP_\star^0}{dP_s^0}(W) \right\} \left\{ \mathbb{E}^0_s \left[D_{\mathrm{UR},s}^{\widehat{\beta}^n,n}(O) \middle| W \right] - \mathbb{E}^0_s \left[D_{\mathrm{UR},s}^{\widehat{\beta}^n,0}(O) \middle| W \right] \right\} \right],
\end{aligned}$$

where the final equality uses that $\mathbb{E}_s^0\left[D_{\mathrm{UR},s}^{\widehat{\beta}^n,0}(O)\middle|W\right]=0$. The second line is $o_P(n_{\min}^{-1/2})$ by C2 and Cauchy-Schwarz. The triangle inequality readily yields that

$$\mathbb{E}_\star^0\left|\mathrm{UR}^{\widehat{\beta}^n,n}(W)-\mathrm{UR}^{\widehat{\beta}^n,0}(W)+\sum_{s=1}^S\mathbb{E}_s^0\left[D_{\mathrm{UR},s}^{\widehat{\beta}^n,n}(O)\middle|W\right]\right|\leq\mathbb{E}_\star^0\left|\mathrm{Rem}_1(\mathcal{P}^n,\mathcal{P}^0)(W)\right|,$$

and the right-hand side is $o_P(n_{\min}^{-1/2})$ by C2. The fact that $\widehat{\omega}^{\widehat{\beta}^n,n}\triangleq\omega^{\widehat{\beta}^n,n}+\sum_{s\in\mathcal{S}}Q_s^n\nabla\Omega_s^{\widehat{\beta}^n,n}$ completes the proof. □

Proof of Theorem A.8. Note that

$$\begin{aligned}\gamma^{\widehat{\beta}^n,0}-\gamma^{\beta^0,0}=&\mathbb{E}_\star^0\left[\left\{\mathrm{UR}^{\widehat{\beta}^n,0}(W)-\mathrm{UR}^{\beta^0,0}(W)\right\}\mathrm{VE}_-^0(W)\right]\\=&\theta^0\left[\omega^{\widehat{\beta}^n,0}-\omega^{\beta^0,0}\right]+\left\{\gamma^{\widehat{\beta}^n,0}-\gamma^{\beta^0,0}-\theta^0\left[\omega^{\widehat{\beta}^n,0}-\omega^{\beta^0,0}\right]\right\}.\end{aligned}$$

By C7, the latter term above is $o_P(n_{\min}^{-1/2})$. By Lemma A.12 and C6,

$$\begin{aligned}\omega^{\widehat{\beta}^n,0}&=\omega^{\widehat{\beta}^n,n}+\sum_{s\in\mathcal{S}}Q_s^n\nabla\Omega_s^{\widehat{\beta}^n,n}-\sum_{s\in\mathcal{S}}(Q_s^n-P_s^0)\nabla\Omega_s^{\widehat{\beta}^n,0}+o_P(n_{\min}^{-1/2})\\&=\mu-\sum_{s\in\mathcal{S}}(Q_s^n-P_s^0)\nabla\Omega_s^{\widehat{\beta}^n,0}+o_P(n_{\min}^{-1/2}).\end{aligned}$$

By C1, $\omega^{\beta^0,0}=\mu$, and so

$$\theta^0\left[\omega^{\widehat{\beta}^n,0}-\omega^{\beta^0,0}\right]=-\theta^0\sum_{s\in\mathcal{S}}(Q_s^n-P_s^0)\nabla\Omega_s^{\widehat{\beta}^n,0}+o_P(n_{\min}^{-1/2}).$$

By C4, we can replace each $\Omega_s^{\widehat{\beta}^n,0}$ above by $\Omega_s^{\beta^0,0}$. □

B.3.3 Terms 2 and 3 under C1'. We first prove Lemma A.9, thereby showing that Term 3 is equal to zero with probability approaching 1. We then establish control over Term 2 by proving Theorem A.10.

Proof of Lemma A.9. By C1', $\theta^0=+\infty$. We show that $\widehat{\omega}^{\beta^0,n}\leq\mu$, and consequently $\widehat{\beta}^n=+\infty$, with probability approaching one. Note that

$$\widehat{\omega}^{\beta^0,n}-\omega^{\beta^0,0}=\sum_{s\in\mathcal{S}}(Q_s^n-P_s^0)\nabla\Omega_s^{\beta^0,n}+\left\{\widehat{\omega}^{\beta^0,n}-\omega^{\beta^0,0}+\sum_{s\in\mathcal{S}}P_s^0\nabla\Omega_s^{\beta^0,n}\right\}.$$

The first term on the right is $O_P(n_{\min}^{-1/2})$ by C3, and the latter term is $o_P(n_{\min}^{-1/2})$ by C2. Hence, $\widehat{\omega}^{\beta^0,n} = \omega^{\beta^0,0} + O_P(n_{\min}^{-1/2})$. By C1', $\omega^{\beta^0,0} < \mu$, and thus $\widehat{\omega}^{\beta^0,n} < \mu$ with probability approaching one. As $\beta^0 \equiv 1$, it must be the case that $\widehat{\theta}^n = +\infty$. □

Proof of Theorem A.10. By Lemma A.9, $\widehat{\theta}^n = +\infty$ with probability approaching 1. Suppose this holds. In this case the TMLE step ensures that $\widehat{\omega}^{\widehat{\beta}^n,n} = Q_\star^n \upsilon^n$. Furthermore, C1' implies that $\omega^{\beta^0,0} = P_\star^0 \upsilon^0$. By the definition of Rem_1, we have that

$$\begin{aligned}\widehat{\omega}^{\widehat{\beta}^n,n} - \omega^{\beta^0,0} &= (Q_\star^n - P_\star^0)\upsilon^n + P_\star^0(\upsilon^n - \upsilon^0) \\ &= (Q_\star^n - P_\star^0)\nabla\Omega_\star^{\widehat{\beta}^n,n} - \sum_{s=1}^S P_\star^0 D_{\mathrm{UR},s}^{\widehat{\beta}^n,n} + P_\star^0 \,\mathrm{Rem}_1(\mathcal{P}^n, \mathcal{P}^0).\end{aligned}$$

By C2, $P_\star^0 \,\mathrm{Rem}_1(\mathcal{P}^n, \mathcal{P}^0) = o_P(n_{\min}^{-1/2})$. Furthermore,

$$\sum_{s=1}^S P_\star^0 D_{\mathrm{UR},s}^{\widehat{\beta}^n,n} = \sum_{s=1}^S P_s^0 \nabla\Omega_s^{\widehat{\beta}^n,n} + \sum_{s=1}^S P_s^0 \left[\frac{dP_\star^0}{dP_s^0} - \frac{dP_\star^n}{dP_s^n}\right] D_{\mathrm{UR},s}^{\widehat{\beta}^n,n}.$$

Using that each $D_{\mathrm{UR},s}^{\beta^0}(O)$ is mean zero for $O \sim P_s^0$ (conditionally on w), the $D_{\mathrm{UR},s}^{\widehat{\beta}^n,n}$ on the right-hand side above can be replaced by $D_{\mathrm{UR},s}^{\widehat{\beta}^n,n} - D_{\mathrm{UR},s}^{\beta^0}$, thereby showing that the latter term above is $o_P(n_{\min}^{-1/2})$ by C2. Hence,

$$\widehat{\omega}^{\widehat{\beta}^n,n} - \omega^{\beta^0,0} = (Q_\star^n - P_\star^0)\nabla\Omega_\star^{\widehat{\beta}^n,n} - \sum_{s=1}^S P_s^0 \nabla\Omega_s^{\widehat{\beta}^n,n} + o_P(n_{\min}^{-1/2}).$$

By Steps 2, 3, and 4 of our estimation procedure and the fact that $\widehat{\theta}^n = +\infty$, $\sum_{s=1}^S Q_s^n \nabla\Omega_s^{\widehat{\beta}^n,n} = 0$ (these steps represent a standard logistic fluctuation submodel for a TMLE, see van der Laan and Rose, 2011). This shows that $\widehat{\omega}^{\widehat{\beta}^n,n} - \omega^{\beta^0,0} = \sum_{s\in\mathcal{S}}(Q_s^0 - P_s^0)\nabla\Omega_s^{\widehat{\beta}^n,n} + o_P(n_{\min}^{-1/2})$. By C4, it follows that $\widehat{\omega}^{\widehat{\beta}^n,n} - \omega^{\beta^0,0} = \sum_{s\in\mathcal{S}}(Q_s^0 - P_s^0)\nabla\Omega_s^{\beta^0,0} + o_P(n_{\min}^{-1/2})$. □

C More interpretable condition for C7

We now provide a more interpretable sufficient condition for C7. First note that Rem_3^n rewrites as

$$\mathrm{Rem}_3^n = \mathbb{E}_\star^0\left[\left\{\mathrm{UR}^{\widehat{\beta}^n,0}(W) - \mathrm{UR}^{\beta^0,0}(W)\right\}\left\{\mathrm{VE}_-^0(W) - \theta^0\right\}\right].$$

The expression in the expectation above is small when either $\widehat{\beta}^n(W) = \beta^0(W)$ or $\mathrm{VE}^0_-(W)$ is close to θ^0. This observation gives some hope that the above expectation will be small when VE^n_- and $\widehat{\theta}^n$ are good estimates of $w \mapsto \mathrm{VE}^0_-(w)$ and θ^0, because it will likely only be most difficult to correctly specify $\mathbb{1}_{\{\mathrm{VE}^n_-(W) \geq \widehat{\theta}^n\}}$ when $\mathrm{VE}^0_-(W) - \theta^0$ is small.

We now make this claim precise. The following margin condition is analogous to that used by Audibert and Tsybakov (2007) in the classification context. In particular, for each $\alpha > 0$ we define Condition α as follows:

$$P_0^\star \left\{ 0 < |\, \mathrm{VE}^0_-(W) - \theta^0| \leq t \right\} \lesssim t^\alpha \text{ for all } t > 0.$$

The above states that $\mathrm{VE}^0_-(W)$ does not place too much mass in the neighborhood of the decision boundary θ^0 that appears in the worst-case unvaccinated risk $\mathrm{UR}^{\beta^0,0}$. The following theorem is an adaptation of Lemma 5.2 in Audibert and Tsybakov (2007). A similar adaptation was given in Luedtke and van der Laan (2016b).

Lemma A.13. *If C1 and Condition α holds for a given $\alpha > 0$, then*

$$|\, \mathrm{Rem}_3^n \,| \lesssim \left\| (\mathrm{VE}^n_-(W) - \widehat{\theta}^n) - (\mathrm{VE}^0_-(W) - \theta^0) \right\|^{2(1+\alpha)/(2+\alpha)}_{2, P_\star^0}$$

$$|\, \mathrm{Rem}_3^n \,| \lesssim \left\| (\mathrm{VE}^n_-(W) - \widehat{\theta}^n) - (\mathrm{VE}^0_-(W) - \theta^0) \right\|^{1+\alpha}_{\infty, P_\star^0}.$$

Shortly we will show that we will show that we can get a faster rate of estimation on the univariate parameter $\widehat{\theta}^n$ than on the infinite-dimensional parameter $w \mapsto \mathrm{VE}^0_-(w)$ when $\alpha \geq 1$. Hence, in this case the above allows us to map our rate of convergence of $w \mapsto \mathrm{VE}^0_-(w)$ into a rate of decay for the remainder term Rem_3^n. Suppose that $\alpha = 1$, which holds if $w \mapsto \mathrm{VE}^0_-(w)$ has bounded Lebesgue density in a neighborhood of θ^0. The supremum norm result in the above lemma suggests that $\mathrm{VE}^n_- - \widehat{\theta}^n$ converging to $\mathrm{VE}^0_- - \theta^0$ at a rate faster than $n_{\min}^{-1/4}$ will suffice to make Rem_3^n negligible.

We close this section by showing the rate of estimation that we can obtain on θ^0 provided the following additional regularity condition is satisfied:

C8) $\left\| \mathrm{VE}^n_- - \mathrm{VE}^0_- \right\|_{2, P_\star^0} = o_P(1)$.

The above is very mild if C2 holds.

Lemma A.14. *If B4, C1, C2, C3, C6, and C8, then $\widehat{\theta}^n = \theta^0 + o_P(1)$. If Condition α also holds, then we have the stronger result that*

$$\widehat{\theta}^n - \theta^0 = O_P\left(n_{\min}^{-1/2} + \left\| \mathrm{VE}^n_- - \mathrm{VE}^0_- \right\|^{\frac{2\alpha}{\alpha+1}}_{2, P_\star^0} \right).$$

Suppose, as is typical, that $\left\|\mathrm{VE}_-^n - \mathrm{VE}_-^0\right\|_{2,P_\star^0}$ shrinks slower than $n_{\min}^{-1/2}$. In this case, if $\alpha \geq 1$, then the above gives conditions under which $\widehat{\theta}^n = \theta^0 + O_P(\left\|\mathrm{VE}_-^n - \mathrm{VE}_-^0\right\|_{2,P_\star^0})$, where the big-oh can be replaced by a little-oh if $\alpha > 1$.

Proof of Lemma A.14. By C6, Lemma A.12, and C3,

$$o_P(n_{\min}^{-1/2}) = \widehat{\omega}^{\widehat{\beta}^n,n} - \mu = \omega^{\widehat{\beta}^n,0} - \omega^{\beta^0,0} + O_P(n_{\min}^{-1/2}).$$

Let $\beta^{n,0}$ denote the function $w \mapsto \mathbb{1}_{\{\mathrm{VE}_-^0(w) < \widehat{\theta}^n\}}$. The above shows that

$$\omega^{\beta^{n,0},0} - \omega^{\beta^0,0} = -[\omega^{\widehat{\beta}^n,0} - \omega^{\beta^{n,0},0}] + O_P(n_{\min}^{-1/2})$$

At the end of this proof, we show that $\omega^{\widehat{\beta}^n,0} - \omega^{\beta^{n,0},0} = o_P(1)$ without using Condition α. For now suppose we have established this. By C1, $\theta \mapsto \omega^{w \mapsto \mathbb{1}\{\mathrm{VE}_-^0(w) < \theta\},0}$ is continuous and increasing at θ^0 so that $\omega^{\beta^{n,0},0} - \omega^{\beta^0,0} = o_P(1)$ is only possible if $\widehat{\theta}^n - \theta^0 = o_P(1)$. If Condition α also holds, then we will show at the end of this proof that $\omega^{\widehat{\beta}^n,0} - \omega^{\beta^{n,0},0} = O_P(\left\|\mathrm{VE}_-^n - \mathrm{VE}_-^0\right\|_{2,P_\star^0}^{\frac{2\alpha}{\alpha+1}})$. Noting that

$$\omega^{\beta^{n,0},0} - \omega^{\beta^0,0} = \frac{\omega^{\beta^{n,0},0} - \omega^{\beta^0,0}}{\widehat{\theta}^n - \theta^0}[\widehat{\theta}^n - \theta^0],$$

we see that C1 and $\widehat{\theta}^n - \theta^0 = o_P(1)$ imply that $\widehat{\theta}^n - \theta^0 = O_P(\omega^{\beta^{n,0},0} - \omega^{\beta^0,0})$, which we have shown to be $O_P(n_{\min}^{-1/2} + \left\|\mathrm{VE}_-^n - \mathrm{VE}_-^0\right\|_{2,P_\star^0}^{\frac{2\alpha}{\alpha+1}})$ under Condition α.

We now establish that $\omega^{\widehat{\beta}^n,0} = \omega^{\beta^{n,0},0} + o_P(1)$ regardless of the validity of Condition α, and that $\omega^{\widehat{\beta}^n,0} = \omega^{\beta^{n,0},0} + O_P(\left\|\mathrm{VE}_-^n - \mathrm{VE}_-^0\right\|_{2,P_\star^0}^{\frac{2\alpha}{\alpha+1}})$ under Condition α. For simplicity we give the proof when $\widehat{\eta}^n = 0$, though the proof for general $\widehat{\eta}^n \in [0,1]$ only differs slightly. Observe that

$$\begin{aligned}
&\left|\omega^{\widehat{\beta}^n,0} - \omega^{\beta^{n,0},0}\right| \\
&= \left|\mathbb{E}_\star^0\left[\{v^0(W) - \lambda^0(W)\}\left\{\mathbb{1}_{\{\mathrm{VE}_-^n(W) < \widehat{\theta}^n\}} - \mathbb{1}_{\{\mathrm{VE}_-^0(W) < \widehat{\theta}^n\}}\right\}\right]\right| \\
&\leq \mathbb{E}_\star^0\left[\left|\mathbb{1}_{\{\mathrm{VE}_-^n(W) < \widehat{\theta}^n\}} - \mathbb{1}_{\{\mathrm{VE}_-^0(W) < \widehat{\theta}^n\}}\right|\right] \\
&\leq P_\star^0\left\{0 \leq |\,\mathrm{VE}_-^0(W) - \widehat{\theta}^n| \leq |\,\mathrm{VE}_-^n(W) - \mathrm{VE}_-^0(W)|\right\}.
\end{aligned}$$

For any $t > 0$, the inequality continues as

$$\leq P_\star^0 \left\{ 0 \leq |\operatorname{VE}_-^0(W) - \widehat{\theta}^n| \leq t \right\} + P_\star^0 \left\{ |\operatorname{VE}_-^n(W) - \operatorname{VE}_-^0(W)| \geq t \right\}$$

$$\leq P_\star^0 \left\{ 0 \leq |\operatorname{VE}_-^0(W) - \widehat{\theta}^n| \leq t \right\} + \frac{\left\| \operatorname{VE}_-^n - \operatorname{VE}_-^0 \right\|_{2,P_\star^0}^2}{t^2}.$$

By C1, $P_\star^0\{\operatorname{VE}_-^0(W) = \widehat{\theta}^n\} = 0$, so that the former term satisfies

$$\lim_{t \downarrow 0} P_\star^0 \left\{ 0 \leq |\operatorname{VE}_-^0(W) - \widehat{\theta}^n| \leq t \right\} = 0.$$

As $\left\| \operatorname{VE}_-^n - \operatorname{VE}_-^0 \right\|_{2,P_\star^0} = o_P(1)$ by C8, one can choose a sequence $\{t_n\}$ and plug it in for t in the preceding inequality for $\left| \omega^{\widehat{\beta}^n,0} - \omega^{\beta^{n,0},0} \right|$ to see that this quantity is $o_P(1)$. If Condition α holds, then one can choose $t_n = \left\| \operatorname{VE}_-^n - \operatorname{VE}_-^0 \right\|_{2,P_\star^0}^{2/(\alpha+1)}$, yielding the stronger result

$$\left| \omega^{\widehat{\beta}^n,0} - \omega^{\beta^{n,0},0} \right| \lesssim \left\| \operatorname{VE}_-^n - \operatorname{VE}_-^0 \right\|_{2,P_\star^0}^{\frac{2\alpha}{\alpha+1}}.$$

□

D Alternative to Steps 2, 3, and 4 in our estimation scheme

We now present alternatives to Steps 2, 3, and 4, to be used when neither (i) $\ell_s(w)$ is a constant multiple of $\sqcup_s(w)$ nor (ii) $\ell_s \equiv 0$ for all s holds. The estimation scheme is identical to that presented in the main text besides the modification of these three steps.

2') Fit a bivariate logistic regression with outcome $(y_s[i] : s = 1, \ldots, S;\ i = 1, \ldots, n_s)$, covariates $\left(\frac{\mathbb{1}_{\{a_s[i]=0\}}\, \ell_s(w)}{n_s P_s^{n,\text{init}}(A=0|w_s[i])} \frac{dP_\star^{n,\text{init}}}{dP_s^{n,\text{init}}}(w_s[i]) : s = 1, \ldots, S;\ i = 1, \ldots, n_s \right)$ and $\left(\frac{\mathbb{1}_{\{a_s[i]=0\}}\, \sqcup_s(w)}{n_s P_s^{n,\text{init}}(A=0|w_s[i])} \frac{dP_\star^{n,\text{init}}}{dP_s^{n,\text{init}}}(w_s[i]) : s = 1, \ldots, S;\ i = 1, \ldots, n_s \right)$, and fixed, subject-level intercept $\left(\operatorname{logit}\left(\mathbb{E}_s^{n,\text{init}}[Y|A=0, w_s[i]] \right) : s = 1, \ldots, S;\ i = 1, \ldots, n_s \right)$. Denote the fitted coefficient in front of the respective covariates by ϵ_n^ℓ and ϵ_n^u.

3') For each $s = 1, \ldots, S$, let $(a, w) \mapsto \mathbb{E}_s^{n,\epsilon_n}[Y|a, w]$ denote the function

$$(a, w) \mapsto \operatorname{logit}^{-1} \Big[\operatorname{logit}\left(\mathbb{E}_s^{n,\text{init}}[Y|a, w] \right)$$

$$+ \epsilon_n^{\ell} \frac{\mathbb{1}_{\{a=0\}} \ell_s(w)}{n_s P_s^{n,\text{init}}(A=0|w)} \frac{dP_\star^{n,\text{init}}}{dP_s^{n,\text{init}}}(w)$$
$$+ \epsilon_n^{u} \frac{\mathbb{1}_{\{a=0\}} \mathsf{u}_s(w)}{n_s P_s^{n,\text{init}}(A=0|w)} \frac{dP_\star^{n,\text{init}}}{dP_s^{n,\text{init}}}(w) \Bigg].$$

4') Let $\mathcal{P}^n = (P_\star^n, P_1^n, \ldots, P_s^n)$ denote any collection of distribution satisfying that, for all (a, w), $\mathbb{E}_s^n[Y|a, w] = \mathbb{E}_s^{n,\epsilon_n}[Y|a, w]$, $P_s^n(A = 1|w) = P_s^{n,\text{init}}(A = 1|w)$, and $\frac{dP_\star^n}{dP_s^n}(w) = \frac{dP_\star^{n,\text{init}}}{dP_s^{n,\text{init}}}(w)$.

E Extensions

E.1 Two-phase sampling

Suppose now that the data is collected via a two-phase sampling scheme in a given trial $s \in \{1, \ldots, S\}$. In particular, suppose that W is collected on only a subset of participants, whereas (L, A, Y) is collected on all participants, where L is a biomarker or collection of biomarkers that may be predictive of W. It is not essential that L happens temporally before A and Y. If no biomarker L is observed, then one can set $L = 0$ for all participants. Let Δ be an indicator of the missingness of W. In this setting, the full data distribution for trial s is $(W, L, A, Y) \sim P_s^{0,F}$, and the observed data structure for trial s is given by $O_s \triangleq (\Delta W, \Delta, L, A, Y) \sim P_s^{0,F}$. We suppose that W is missing at random, in the sense that $\Delta \perp\!\!\!\perp W|(L, A, Y)$ for each trial s, and further that censoring mechanism, i.e. the probability that $\Delta = 1$ given each realization of (L, A, Y), is known. This ensures that $P_s^{0,F}$ can be identified with P_s^0 by the G-computation formula Robins (1986). In particular, for any event E on O_s, we have that $P_s^{0,F}\{E\} = \mathbb{E}_s^0[P_s^0\{E|\Delta = 1, L\}]$.

A common sampling scheme that generates such data in vaccine efficacy trials is a nested case-control sampling scheme (Breslow, 1996), where the outcome W is observed on all cases (participants with $Y = 1$), and the outcome W is only observed on a subset of controls ($Y = 0$). Often these sampling schemes will take the form of an m:1 scheme, such that, for each case with W observed, W is observed for m controls. While technically the indicator Δ is drawn without replacement for these m individuals, one can typically ignore this dependence in the data with little impact on precision or coverage.

A simple modification of our procedure via inverse probability weighting allows estimation of our efficacy lower bound ϕ^0. For efficiency gains, we recommend estimating the censoring mechanism even though it is known (see

Theorem 2.3 in van der Laan and Robins, 2003). While estimating the censoring mechanism improves the precision of our estimator, it does not reduce the width of our confidence intervals. This thus leads to a conservative inferential procedure. The proof of correctness of this approach is beyond the scope of this work, though closely follows the arguments given in Rose and van der Laan (2011), with minor tweaks to account for the multiple sample nature of the problem. Indeed, Steps 2, 3, and 4 constitute an IPCW-TMLE, as presented in Rose and van der Laan (2011). For simplicity we only give the algorithm for the case presented in Section 4.1.1 in the main text, namely where either (i) the chosen $\ell_s(w)$ is a constant multiple of the chosen $\mathsf{u}_s(w)$, where the multiple is independent of s and w, or (ii) the chosen $\ell_s \equiv 0$ for all s so that $\lambda^0 \equiv 0$. The modification of Steps 2, 3, and 4 is analogous to the modifications made to the algorithm in the main text given in Appendix D. We start at Step 0 so that the other steps parallel those given for the algorithm in the main text.

To emphasize the fact that many of the parameters below depend on the full data structure (W, L, A, Y) rather than the (censored) observed data structure $(\Delta W, \Delta, L, A, Y)$, we replace the "0" in the superscript by "$0, F$" when the parameter is defined for the full data distribution, e.g. we write $\lambda^{0,F}$ rather than λ^0. Similarly, we write $\lambda^{n,F}$ rather than λ^n when denoting estimates of parameters of the full data distribution. Finally, we note that below we denote the observed value of L and Δ for participant i from trial s by $l_s[i]$ and $\delta_s[i]$, respectively.

0. Estimate each $(l, a, y) \mapsto P_s^0\{\Delta = 1|l, a, y\}$ using a completed trial s-specific correctly specified parametric model. Standardize these estimates by an appropriate constant so that, for each $s = 1, \ldots, S$,

$$\sum_{i=1}^{n_s} \frac{\delta_s[i]}{P_s^n\{\Delta = 1|l_s[i], a_s[i], y_s[i]\}} = n_s.$$

   Note: correct parametric model specification is possible in this context due to the presumed knowledge (by experimental design) of the censoring mechanism $P_s^0\{\Delta = 1|l, a, y\}$.

1. Let $(a, w) \mapsto \mathbb{E}_s^{n,F,\text{init}}[Y|a, w]$, $w \mapsto P_s^{n,F,\text{init}}(A = 1|w)$, and $w \mapsto \frac{dP_\star^{n,F,\text{init}}}{dP_s^{n,F,\text{init}}}(w)$ represent estimates of $(a, w) \mapsto \mathbb{E}_s^{0,F}[Y|a, w]$, $w \mapsto P_s^{0,F}(A = 1|w)$, and $w \mapsto \frac{dP_\star^0}{dP_s^{0,F}}(w)$, respectively.
   Note: Rose and van der Laan (2011) describe a weighted loss-based estimation scheme that leverages the information in the biomarker L when

estimating these quantities. This procedure makes use of the estimate of the censoring mechanism from Step 0. Rose and van der Laan (2011) also extend the super-learner of van der Laan et al. (2007) to two-phase sampling designs.

2. Fit a weighted univariate logistic regression with weights $\left(\frac{\Delta}{P_s^n\{\delta_s[i]=1|l_s[i],a_s[i],y_s[i]\}} : s=1,\ldots,S;\ i=1,\ldots,n_s\right)$, outcome $(y_s[i] : s=1,\ldots,S;\ i=1,\ldots,n_s)$, covariate $\left(\frac{\mathbb{1}_{\{a_s[i]=0\}}\,\mathsf{u}_s(w)}{n_s P_s^{n,F,\text{init}}(A=0|w_s[i])}\frac{dP_\star^{n,F,\text{init}}}{dP_s^{n,F,\text{init}}}(w_s[i]) : s=1,\ldots,S;\ i=1,\ldots,n_s\right)$, fixed, subject-level intercept $\left(\text{logit}\left(\mathbb{E}_s^{n,F,\text{init}}[Y|A=0,w_s[i]]\right) : s=1,\ldots,S;\ i=1,\ldots,n_s\right)$. Denote the fitted coefficient in front of the covariate by ϵ_n.

3. For each $s=1,\ldots,S$, let $(a,w)\mapsto \mathbb{E}_s^{n,F,\epsilon_n}[Y|a,w]$ denote the function

$$(a,w)\mapsto \text{logit}^{-1}\left[\text{logit}\left(\mathbb{E}_s^{n,F,\text{init}}[Y|a,w]\right)+\epsilon_n\frac{\mathbb{1}_{\{a=0\}}\,\mathsf{u}_s(w)}{n_s P_s^{n,F,\text{init}}(A=0|w)}\frac{dP_\star^{n,F,\text{init}}}{dP_s^{n,F,\text{init}}}(w)\right].$$

4. Let $\mathcal{P}^{n,F}=(P_\star^n,P_1^{n,F},\ldots,P_s^{n,F})$ denote any collection of distribution satisfying that, for all (a,w), $\mathbb{E}_s^{n,F}[Y|a,w]=\mathbb{E}_s^{n,F,\epsilon_n}[Y|a,w]$, $P_s^{n,F}(A=1|w)=P_s^{n,F,\text{init}}(A=1|w)$, and $\frac{dP_\star^n}{dP_s^{n,F}}(w)=\frac{dP_\star^{n,F,\text{init}}}{dP_s^{n,F,\text{init}}}(w)$. Furthermore, for each completed trial s, let $Q_s^{n,F}$ denote the distribution that puts mass proportional to $\delta_s[i]/(n_s P_s^n\{\Delta=1|l_s[i],a_s[i],y_s[i]\})$ at each observation $i=1,\ldots,n_s$, and zero mass elsewhere.

5. For each $\beta:\mathcal{W}\to[0,1]$, let $\widehat{\omega}^{\beta,n,F}\triangleq\omega^{\beta,n,F}+\sum_{s\in\mathcal{S}}Q_s^{n,F}\nabla\Omega_s^{\beta,n,F}$, and note that $\widehat{\omega}^{\beta,n,F}$ rewrites as $Q_\star^{n,F}\,\text{UR}^{\beta,n,F}+\sum_{s=1}^S Q_s^{n,F}\nabla\Omega_s^{\beta,n,F}$.

6. Let $\widehat{\theta}^{n,F}\triangleq\sup\{\theta:\widehat{\omega}^{w\mapsto\mathbb{1}_{\{\text{VE}_-^{n,F}(w)<\theta\}},n,F}\leq\mu\}$, where $\sup\emptyset=-\infty$.

7. Let $\widehat{\eta}^{n,F}$ be any element of the set $\text{argmin}_{\eta\in[0,1]}\left(\widehat{\omega}^{\beta_\eta,n,F}-\mu\right)^2$, where $\beta_\eta\triangleq w\mapsto\mathbb{1}_{\{\text{VE}_-^{n,F}(w)<\widehat{\theta}^{n,F}\}}+\eta\,\mathbb{1}_{\{\text{VE}_-^{n,F}(w)=\widehat{\theta}^{n,F}\}}$.

8. Let $\widehat{\beta}^{n,F}\triangleq\beta_{\widehat{\eta}^{n,F}}$.

9. Estimate $\gamma^{\widehat{\beta}^{n,F},0}$ with

$$\begin{aligned}\widehat{\gamma}^{\widehat{\beta}^{n,F},n,F}&\triangleq\gamma^{\widehat{\beta}^{n,F},n,F}+\sum_{s\in\mathcal{S}}Q_s^{n,F}\nabla\Gamma_s^{\widehat{\beta}^{n,F},n,F}\\&=n_\star^{-1}\sum_{i=1}^{n_\star}\text{UR}^{\widehat{\beta}^{n,F},n,F}(w_\star[i])\,\text{VE}^{n,F}(w_\star[i])+\sum_{s=1}^S Q_s^{n,F}\nabla\Gamma_s^{\widehat{\beta}^{n,F},n,F}.\end{aligned}$$

10. Estimate ϕ^0 with $\widehat{\phi}^{n,F} \triangleq \frac{\widehat{\gamma}^{\widehat{\beta}^{n,F},n,F}}{\widehat{\omega}^{\widehat{\beta}^{n,F},n,F}}$.

For confidence interval construction, we note that similar conditions to those used in 3 yield that our estimator is asymptotically linear, with the same $P_\star^0$ gradient as in the main text (though with the observed data parameters for the completed trial distributions replaced by full data parameters) and, for $s = 1, \ldots, S$, P_s^0 gradients $o_s \mapsto \frac{\delta}{P_s^0\{\Delta=1|l,a,y\}} \nabla \Phi_s^0(o_s)$, again replacing the observed data parameters by the full data parameters in the definition of $\nabla \Phi_s^0(o_s)$.

Remark 6. Consider a randomized trial where L precedes randomization and A is independent of L conditional on W. In this case, it is straightforward to improve the efficiency of the above procedure if L is predictive of Y after accounting for A and W in at least one of the trials. In particular, efficiency could be improved by leveraging this biomarker when estimating both the (known) propensity score $P_s^{0,F}\{A|W\}$ and the outcome regression $\mathbb{E}_s^{0,F}[Y|A,W]$, namely by adding L to the both of the corresponding conditioning statements. The same efficiency gain of course holds for the algorithm in the main text, since setting $\Delta = 1$ with probability one shows that W being observed on all individuals is a special case of the results in this appendix. □

E.2 Monotonic vaccine efficacy curve

In this section, we describe a situation in which one can replace the condition that $P_\star^0\{\mathrm{VE}_-^0(W) = \theta^0\} = 0$ with the following new condition:

M) W is real-valued and $w \mapsto \mathrm{VE}_-^0(w)$ is monotonic.

We break this section into three parts. First, we present a new partial bridging formula specific to the monotonic VE_-^0 case. Then, we give formal conditions that will allow one to establish the pathwise derivative of the parameter specified by this formula. Finally, we describe how to modify our estimator from the main text so that the validity of the confidence intervals neither relies on $P_\star^0\{\mathrm{VE}_-^0(W) = \theta^0\} = 0$ nor on C7.

E.2.1 Partial bridging formula. We will use the notation from the main text to express our partial bridging formula. We will define alternatives to θ^0, η^0, β^0, and ϕ^0, which we will respectively denote by $\underline{\theta}^0$, $\underline{\eta}^0$, $\underline{\beta}^0$, and $\underline{\phi}^0$.

Define

$$\underline{\theta}^0 \triangleq \sup\left\{\theta \in \mathbb{R} : \omega^{w \mapsto \mathbb{1}\{w<\theta\},0} \leq \mu\right\},$$

where $\sup \emptyset = -\infty$ by convention. Let $\underline{\beta}_{\underline{\eta}} \triangleq w \mapsto \mathbb{1}_{\{w<\underline{\theta}^0\}} + \underline{\eta}\, \mathbb{1}_{\{w=\underline{\theta}^0\}}$, and define $\underline{\eta}^0$ to be the smallest element of the set $\operatorname{argmin}_{\eta\in[0,1]} \left(\omega^{\underline{\beta}_\eta,0} - \mu\right)^2$. Let $\underline{\beta}^0 \triangleq \underline{\beta}_{\underline{\eta}^0}$. One can show that, if $\mathbb{E}^0_\star[\lambda^0(W)] \leq \mu \leq \mathbb{E}^0_\star[\upsilon^0(W)]$, then $\omega^{\underline{\beta}^0,0} = \mu$. Our partial bridging parameter is given by

$$\underline{\Phi}(\mathcal{P}^0) \triangleq \frac{\gamma^{\underline{\beta}^0,0}}{\omega^{\underline{\beta}^0,0}} \triangleq \underline{\phi}^0.$$

The proof of the following lemma is nearly identical to the proof of Lemma 1 and so is omitted.

Lemma A.15. *Suppose M, B1, B2, B3, and B4 hold. If* $\mu = \mathbb{E}^0_\star[\mathbb{E}^{0,F}_\star[Y|A = 0, W]]$, *then* $\Psi(P^{0,F}_\star) \geq \underline{\phi}^0$.

E.2.2 First-order expansion of $\underline{\phi}^0$. Our first-order expansion also replaces C1 in the main text with two alternative assumptions, the first of which is given below.

D1) $\mathbb{E}^0_\star[\lambda^0(W)] < \mu < \mathbb{E}^0_\star[\upsilon^0(W)]$ and either (i) $P^0_\star\{W = \theta^0\} > 0$ or (ii) $P^0_\star\{W = \theta^0\} = 0$, $w \mapsto \mathrm{VE}^0_-(w)$ is Lipschitz at θ^0, and

$$0 < \liminf_{t\to 0} \frac{P^0_\star\{W < \theta + t\} - P^0_\star\{W < \theta\}}{t}.$$

Note that $\mathbb{E}^0_\star[\lambda^0(W)] < \mu < \mathbb{E}^0_\star[\upsilon^0(W)]$ implies that θ^0 is finite.

We give a theorem presenting the gradients of our parameter. For brevity, we only sketch the proof.

Theorem A.16. *If B3 and B4 hold and either D1, C1', or C1" holds, then* $\underline{\Phi}$ *is pathwise differentiable and, for each* $s \in \mathcal{S}$, *the* P^0_s *gradient is given by*

$$\nabla\underline{\Phi}^0_s(o_s) = \begin{cases} \dfrac{\nabla\Gamma^{\underline{\beta}^0,0}_s(o_s)}{\omega^{\underline{\beta}^0,0}} - \mathrm{VE}^0_-(\underline{\theta}^0)\dfrac{\nabla\Omega^{\underline{\beta}^0,0}_s(o_s)}{\omega^{\underline{\beta}^0,0}}, & \text{if } \mathbb{E}^0_\star[\lambda^0(W)] < \mu < \mathbb{E}^0_\star[\upsilon^0(W)], \\ \dfrac{\nabla\Gamma^{\underline{\beta}^0,0}_s(o_s)}{\omega^{\underline{\beta}^0,0}} - \gamma^{\underline{\beta}^0,0}\dfrac{\nabla\Omega^{\underline{\beta}^0,0}_s(o_s)}{[\omega^{\underline{\beta}^0,0}]^2}, & \text{otherwise.} \end{cases}.$$

Sketch of Proof of Theorem A.16. The proof under C1' or C1" is essentially identical to that of Theorem 2 under the same conditions, we suppose D1 in the remainder.

We first outline the convergence result of $\underline{\beta}^\epsilon$ to $\underline{\beta}^0$. If (i), then we can instead show that $\underline{\theta}^\epsilon = \underline{\theta}^0$ for all ϵ small enough and $\underline{\eta}^\epsilon = \underline{\eta}^0 + O(\epsilon)$. If (ii), then D1 yields that $\underline{\theta}^\epsilon = \underline{\theta}^0 + O(\epsilon)$. In either case, $\underline{\theta}^0$ is finite for all ϵ small enough so that $\omega^{\underline{\beta}^\epsilon,\epsilon} = \mu$ for all ϵ small enough. Hence,

$$\frac{\phi^\epsilon - \phi^0}{\epsilon} = \mu^{-1}\frac{\gamma^{\underline{\beta}^\epsilon,\epsilon} - \gamma^{\underline{\beta}^0,0}}{\epsilon},$$

so it is enough to study the right-hand side multiplied by $\mu\epsilon$. We will use that

$$\begin{aligned}\gamma^{\underline{\beta}^\epsilon,\epsilon} - \gamma^{\underline{\beta}^0,0} = &\left\{\gamma^{\underline{\beta}^0,\epsilon} - \gamma^{\underline{\beta}^0,0} - \mathrm{VE}^0_-(\underline{\theta}^0)[\omega^{\underline{\beta}^0,\epsilon} - \omega^{\underline{\beta}^0,0}]\right\} \\ &+ \mathrm{VE}^0_-(\underline{\theta}^0)\left[\omega^{\underline{\beta}^\epsilon,\epsilon} - \omega^{\underline{\beta}^0,0}\right] \\ &+ \left\{\gamma^{\underline{\beta}^\epsilon,\epsilon} - \gamma^{\underline{\beta}^0,\epsilon} - \mathrm{VE}^0_-(\underline{\theta}^0)[\omega^{\underline{\beta}^\epsilon,\epsilon} - \omega^{\underline{\beta}^0,\epsilon}]\right\}\end{aligned}$$

A delta method argument shows that dividing the leading term by ϵ and taking the limit as $\epsilon \to 0$ yields $\sum_{s\in\mathcal{S}} P^0_s h_s[\nabla\Gamma^{\underline{\beta}^0,0}_s(o_s) + \mathrm{VE}^0_-(\underline{\theta}^0)\nabla\Omega^{\underline{\beta}^0,0}_s]$. As we established earlier in this proof, $\omega^{\underline{\beta}^\epsilon,\epsilon} = \mu = \omega^{\underline{\beta}^0,0}$ for all ϵ small enough, so that the second term above is zero for all ϵ small enough . The remainder of this proof aims to show that the final term is $o(\epsilon)$. We will use a constant c that may vary line by line. We now study the final term above, whose numerator bounds as

$$\begin{aligned}&\left|\gamma^{\underline{\beta}^\epsilon,\epsilon} - \gamma^{\underline{\beta}^0,\epsilon} - \mathrm{VE}^0_-(\underline{\theta}^0)[\omega^{\underline{\beta}^\epsilon,\epsilon} - \omega^{\underline{\beta}^0,\epsilon}]\right| \\ &= (1 + c|\epsilon|)\,\mathbb{E}^0_\star\left|\{\upsilon^\epsilon(W) - \lambda^\epsilon(W)\}\{\underline{\beta}^\epsilon(W) - \underline{\beta}^0(W)\}\left\{\mathrm{VE}^\epsilon_-(W) - \mathrm{VE}^0_-(\underline{\theta}^0)\right\}\right| \\ &\leq (1 + c|\epsilon|)\,\mathbb{E}^0_\star\left[|\underline{\beta}^\epsilon(W) - \underline{\beta}^0(W)|\left\{\left|\mathrm{VE}^0_-(W) - \mathrm{VE}^0_-(\underline{\theta}^0)\right| + c|\epsilon|\right\}\right].\end{aligned}$$

Call the right-hand side $\underline{T}^\epsilon_1$. If (i) holds, then the fact that $\underline{\theta}^\epsilon = \underline{\theta}^0$ for all ϵ small enough shows that, for all such ϵ,

$$\underline{T}^\epsilon_1 \leq c|\epsilon|(1 + c|\epsilon|)|\underline{\eta}^\epsilon - \underline{\eta}^0|P^0_\star\{W = \underline{\theta}^0\}.$$

As $\underline{\eta}^\epsilon - \underline{\eta}^0 = O(\epsilon)$, the right-hand side must be $o(\epsilon)$. If, instead, (ii) holds, then similar techniques to those used to control T^ϵ_1 in Theorem A.3 show that

$$\underline{T}^\epsilon_1 \leq (1 + c|\epsilon|)\,\mathbb{E}^0_\star\left[\mathbb{1}_{\{0<|W-\underline{\theta}^0|<|\underline{\theta}^\epsilon-\underline{\theta}^0|\}}\left\{\left|\mathrm{VE}^0_-(W) - \mathrm{VE}^0_-(\underline{\theta}^0)\right| + c|\epsilon|\right\}\right].$$

By the fact that VE^0_- is Lipschitz at $\underline{\theta}^0$ and the bound on $|W - \underline{\theta}^0|$ given in the indicator, we have that

$$\underline{T}^\epsilon_1 \leq (1 + c|\epsilon|)\left\{c\left|\underline{\theta}^\epsilon - \underline{\theta}^0\right| + c|\epsilon|\right\}P^0_\star\left\{0 < |W - \underline{\theta}^0| < |\underline{\theta}^\epsilon - \underline{\theta}^0|\right\}.$$

As $\underline{\theta}^\epsilon - \underline{\theta}^0 = O(\epsilon)$ and the probability statement is $o(1)$, the right-hand side is $o(\epsilon)$. □

E.2.3 Modification to our estimator. We now present the modification to the estimator presented in the main text that allows us to replace the condition that $\mathrm{VE}^0_-(W)$, $W \sim P^0_\star$, does not concentrate mass at a decision boundary with the condition that $W \sim P^0_\star$ does not concentrate mass at a decision boundary and the monotonicity condition M.

The first 5 steps of our estimation procedure are identical to that in the main text, and the modification of the remaining steps is given below.

6. Let $\widehat{\underline{\theta}}^n \triangleq \sup\{\theta : \widehat{\omega}^{w \mapsto \mathbb{1}_{\{w<\theta\}},n} \leq \mu\}$, where $\sup \emptyset = -\infty$.

7. Let $\widehat{\underline{\eta}}^n$ be any element of the set $\operatorname{argmin}_{\eta \in [0,1]} \left(\widehat{\omega}^{\beta_\eta,n} - \mu\right)^2$, where $\beta_\eta \triangleq w \mapsto \mathbb{1}_{\{w<\widehat{\underline{\theta}}^n\}} + \eta\, \mathbb{1}_{\{w=\widehat{\underline{\theta}}^n\}}$.

8. Let $\widehat{\underline{\beta}}^n \triangleq \beta_{\widehat{\underline{\eta}}^n}$.

9. Estimate $\gamma^{\widehat{\underline{\beta}}^n,0}$ with

$$\begin{aligned}\widehat{\gamma}^{\widehat{\underline{\beta}}^n,n} &\triangleq \gamma^{\widehat{\underline{\beta}}^n,n} + \sum_{s \in \mathcal{S}} Q^n_s \nabla \Gamma^{\widehat{\underline{\beta}}^n,n}_{\underline{s}} \\ &= n_\star^{-1} \sum_{i=1}^{n_\star} \mathrm{UR}^{\widehat{\underline{\beta}}^n,n}_-(w_\star[i])\, \mathrm{VE}^n(w_\star[i]) + \sum_{s=1}^{S} Q^n_s \nabla \Gamma^{\widehat{\underline{\beta}}^n,n}_{\underline{s}}.\end{aligned}$$

10. Estimate $\underline{\phi}^0$ with $\widehat{\underline{\phi}}^n \triangleq \frac{\widehat{\gamma}^{\widehat{\underline{\beta}}^n,n}}{\widehat{\omega}^{\widehat{\underline{\beta}}^n,n}}$.

The conditions for the asymptotic linearity of this approach are similar to those of Theorem 3, but weaker because they do not rely on the condition that $P^0_\star\{\mathrm{VE}^0_-(W) = \theta^0\} = 0$ (nor any analogue thereof). To start, we consider why, at least for W continuous, $P^0_\star\{\mathrm{VE}^0_-(W) = \theta^0\} = 0$ was essentially necessary for it to be plausible that the procedure in the main text was asymptotically linear. For asymptotic linearity to hold, the indicators that $\mathrm{VE}^n_-(w) < \widehat{\theta}^n$ and $\mathrm{VE}^n_-(w) = \widehat{\theta}^n$ that appear in the empirical gradients must converge to a fixed limit. The fact that VE^n_- is only an estimate of VE^0_- suggests that it is not plausible that the indicators of the events $\mathrm{VE}^n_-(w) < \widehat{\theta}^n$ or $\mathrm{VE}^n_-(w) = \widehat{\theta}^n$ converge to a fixed limit for any w for which $\mathrm{VE}^0_-(w) = \theta^0$. Thus, if W is continuous and $P^0_\star\{\mathrm{VE}^0_-(W) = \theta^0\} > 0$, then it is not plausible that these indicators converge to any fixed limit.

Consider now the procedure above. The procedure above replaces the indicators that $\mathrm{VE}^n_-(w) < \widehat{\theta}^n$ and $\mathrm{VE}^n_-(w) = \widehat{\theta}^n$ with indicators that $w < \widehat{\underline{\theta}}^n$

and $w = \widehat{\underline{\theta}}^n$: in particular, it is expected to be valid regardless of the value of $P_\star^0\{W = \underline{\theta}^0\}$. The conditions needed to ensure that these new indicators converge to a fixed limit are weaker than those needed for the convergence of $\mathrm{VE}_-^n(w) < \widehat{\theta}^n$ or $\mathrm{VE}_-^n(w) = \widehat{\theta}^n$. If $P_\star^0$ concentrates mass at $\underline{\theta}^0$, then we expect that $\widehat{\underline{\theta}}^n$ will equal $\underline{\theta}^0$ with probability approaching one, so that clearly the indicators that $w < \widehat{\underline{\theta}}^n$ and $w = \widehat{\underline{\theta}}^n$ are fixed with probability approaching one. If, on the other hand, $P_\star^0$ does not concentrate mass at $\underline{\theta}^0$, then indicators that $w < \widehat{\underline{\theta}}^n$ and $w = \widehat{\underline{\theta}}^n$ will generally converge to a fixed limit if $\widehat{\underline{\theta}}^n \to \underline{\theta}^0$. For some intuition on why $\widehat{\underline{\theta}}^n \to \underline{\theta}^0$ is to be expected, note that $\widehat{\underline{\theta}}^n$ is essentially an empirical quantile for an $(\upsilon^0 - \lambda^0)$-weighted version of $P_\star^0$. We say "essentially" because conditions can be given under which we can replace the estimated weights $\upsilon^n - \lambda^n$ by the true weights $\upsilon^0 - \lambda^0$ at the expense of an $O_P(n_{\min}^{-1/2})$ term. Given this replacement, we can show that D1 yields that $\widehat{\underline{\theta}}^n = \theta^0 + O_P(n_{\min}^{-1/2})$. Note that this is stronger than the slower than root-n rate of convergence of $\widehat{\theta}^n$ to θ^0 in the nonmonotonic case given in Lemma A.14.

F Nested case-control simulation results

Figure 3 shows that our the coverage of our estimation scheme decreased under the 1:1 nested case-control design, though also that this coverage improves with sample size. At the Moderate and Tight settings, Figure 4 shows that our estimate of the partially bridged lower bound has positive bias for the nested case-control design, though that this bias reduces with sample size.

Figure 5 compares the Monte Carlo variance of the estimator that uses the known censoring mechanism versus the estimator that estimates the censoring mechanism. The two procedures appear to have similar variance in this particular simulation setting.

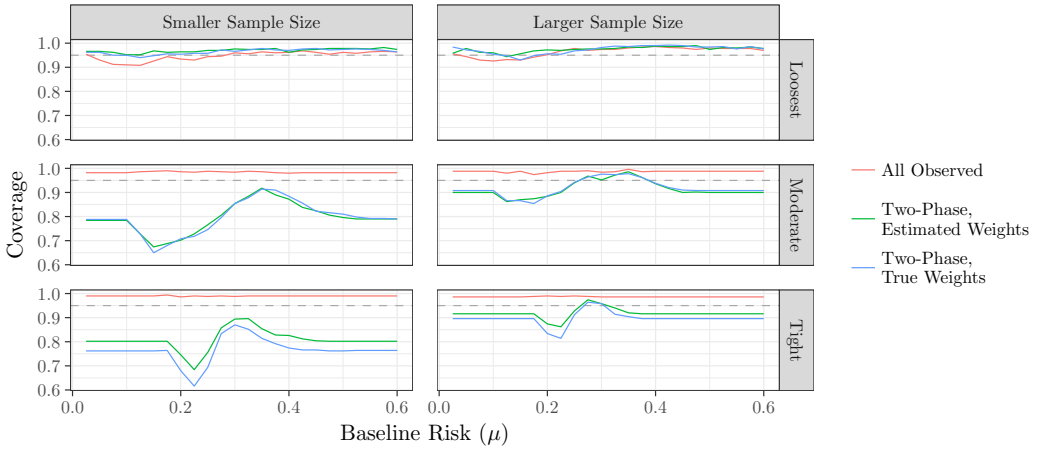

Figure 3: Comparison of the coverage of our lower confidence bound for ϕ^0, i.e. the lower bound on the vaccine efficacy, when there is and is not data missing due to a nested case-control sampling design. Conducted at both smaller and larger sample sizes, respectively with $(n_\star, n_1, n_2)$ equal to $(100, 2000, 2000)$ and $(200, 4000, 4000)$, and for different choices of ℓ_s and u_s, determining the tightness of the unvaccinated risk bounds. Horizontal dashed lines drawn at 95% coverage, vertical dashed lines drawn at the true baseline risk value μ.

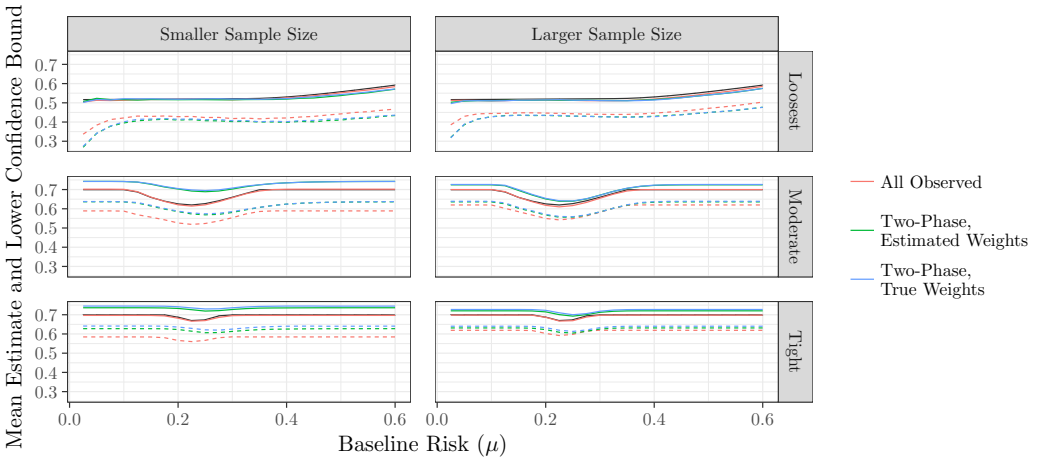

Figure 4: Comparison of average estimates (solid lines) and lower confidence bounds (dashed lines) for ϕ^0, i.e. the lower bound on the vaccine efficacy, when there is and is not data missing due to a nested case-control sampling design. Conducted at both smaller and larger sample sizes, respectively with $(n_\star, n_1, n_2)$ equal to $(100, 2000, 2000)$ and $(200, 4000, 4000)$, and for different choices of ℓ_s and u_s, determining the tightness of the unvaccinated risk bounds. Black trend lines denote true $(\ell_s, \mathsf{u}_s, \mu)$-specific lower bound.

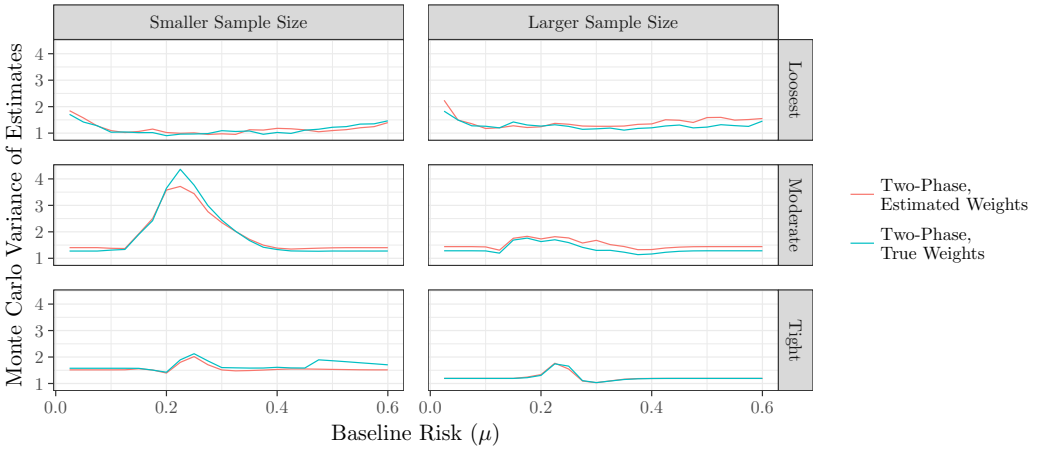

Figure 5: Comparison of estimator's Monte Carlo variance in the presence of two-phase sampling (standardized by Monte Carlo variance of estimator that observed W for all individuals). Conducted at both smaller and larger sample sizes, respectively with $(n_\star, n_1, n_2)$ equal to $(100, 2000, 2000)$ and $(200, 4000, 4000)$, and for different choices of ℓ_s and u_s, determining the tightness of the unvaccinated risk bounds.